



**Actuarial
Research Centre**

Institute and Faculty
of Actuaries

Actuarial Research Centre (ARC)

PhD studentship output

The Actuarial Research Centre (ARC) is the Institute and Faculty of Actuaries' network of actuarial researchers around the world.


The ARC seeks to deliver research programmes that bridge academic rigour with practitioner needs by working collaboratively with academics, industry and other actuarial bodies.

The ARC supports actuarial researchers around the world in the delivery of cutting-edge research programmes that aim to address some of the significant challenges in actuarial science.

EMPIRICAL STUDIES IN CORPORATE CREDIT
MODELLING; LIQUIDITY PREMIA, FACTOR
PORTFOLIOS & MODEL UNCERTAINTY

by

Paul René Frank van Loon



Submitted for the degree of
Doctor of Philosophy

DEPARTMENT OF ACTUARIAL MATHEMATICS AND STATISTICS
SCHOOL OF MATHEMATICAL AND COMPUTER SCIENCES
HERIOT-WATT UNIVERSITY

March 2017

The copyright in this thesis is owned by the author. Any quotation from the report or use of any of the information contained in it must acknowledge this report as the source of the quotation or information.

Abstract

Insurers match the cash flows of typically illiquid insurance liabilities, such as in-force annuities, with government and corporate bonds. As they intend to buy corporate bonds and hold them to maturity, they can capture the value attached to liquidity, without running the market liquidity risk that is associated with having to sell bonds in the open market. During the long consultation period dedicated to the mark-to-market valuation of insurance assets and liabilities for the Solvency II regulatory framework, CEIOPS noted the importance of the accurate breakdown of the credit spread into its components, most notably the credit and non-credit (i.e. liquidity) components. In this thesis we review many modelling efforts to isolate the liquidity premium and propose a reduced-form modelling approach that relies on a new, relative liquidity proxy.

Challenging the status quo when it comes to active and passive investment strategies, products and funds, Exchange Traded Funds and ‘smart-beta’ products provide investors with straightforward ways to strategically expose a portfolio to risk drivers, raising the bar for traditional investment funds and managers. In this thesis, we investigate how traditional sources of equity outperformance (alpha), such as small caps, low volatility and value, translate to UK corporate bonds. For automated trading strategies in corporate bonds, and those with specific factor exposure requirements in particular, transaction costs, rebalancing and an optimal turnover strategy are crucial; these aspects of building factor portfolios are explored for the UK market.

Since the financial crisis, mathematical models used in finance have been subject to a fair amount of criticism. More than ever has this highlighted the need of better risk management of financial models themselves, leading to a surge in ‘model validation’ roles in industry and an increased scrutiny from regulatory bodies. In this thesis we look at stochastic credit models that are commonly used by insurers to project forward credit-risky bond portfolios and the model uncertainty and parameter risk that arises as a result of relying on published credit migration matrices. Specifically, our investigation focuses on two violations of the Markovian process that credit transitions are assumed to follow and statistical uncertainty of the migration matrix.

Acknowledgements

First of all, I would like to thank Professor Andrew Cairns and Professor Alexander McNeil for providing guidance on the research conducted during my time as a PhD candidate. Regular discussions with both Andrew and Alex have been instrumental over the last three and a half years. Their expertise has been of great benefit in aiding my understanding of some technical aspects of my studies and their enthusiasm helped shape the direction of my research. Their comments on the initial draft contributed to an improved final thesis.

Secondly, I would like to thank Alex Veys from Partnership Assurance. Not only have Partnership supported my PhD studentship financially, regular meetings with Alex during the first year of the PhD deepened my understanding of corporate bond markets and shaped the direction of the research going forward. I am grateful for the time I got to spend with Alex and the rest of the investment team at Partnership's offices during the summer of 2014. During this time I was fortunate enough to have the opportunity to discuss my progress with many industry experts including Etienne Comon from GS Asset Management, Paul Sweeting from JPM Asset Management, Albert Desclee from Barclays Capital and others.

I also wish to extend my gratitude to the Institute and Faculty of Actuaries and the Actuarial Research Centre for their support throughout my PhD. In particular I would like to thank Kevin and Sarah from the IFoA, who have been ever so helpful in accommodating conference speaking opportunities in Washington, Kuala Lumpur, Beijing, Dublin and many places around the UK. I thank Con Keating for inviting me to present some of my work to the CISI Bond Group in September 2014.

Finally, I would like to thank my family for their support. I am grateful to Marijke, Eric, Megan and Mark for their constant love, support and encouragement throughout the years.

ACADEMIC REGISTRY
Research Thesis Submission



Name:	Paul Rene Frank van Loon		
School/PGI:	Mathematics and Computer Science		
Version: <i>(i.e. First, Resubmission, Final)</i>	Final	Degree Sought (Award and Subject area)	Doctor of Philosophy

Declaration

In accordance with the appropriate regulations I hereby submit my thesis and I declare that:

- 1) the thesis embodies the results of my own work and has been composed by myself
- 2) where appropriate, I have made acknowledgement of the work of others and have made reference to work carried out in collaboration with other persons
- 3) the thesis is the correct version of the thesis for submission and is the same version as any electronic versions submitted*.
- 4) my thesis for the award referred to, deposited in the Heriot-Watt University Library, should be made available for loan or photocopying and be available via the Institutional Repository, subject to such conditions as the Librarian may require
- 5) I understand that as a student of the University I am required to abide by the Regulations of the University and to conform to its discipline.

* Please note that it is the responsibility of the candidate to ensure that the correct version of the thesis is submitted.

Signature of Candidate:		Date:	
-------------------------	--	-------	--

Submission

Submitted By <i>(name in capitals)</i> :	
Signature of Individual Submitting:	
Date Submitted:	

For Completion in the Student Service Centre (SSC)

Received in the SSC by <i>(name in capitals)</i> :			
<i>Method of Submission</i> <i>(Handed in to SSC; posted through internal/external mail):</i>			
<i>E-thesis Submitted (mandatory for final theses)</i>			
Signature:		Date:	

Contents

1	Motivation to the PhD Thesis	1
1.1	Regulatory Environment of the Liquidity Premium in Solvency II . . .	2
1.2	Rise of ETFs and Smart Beta Products	6
1.3	Financial Risk of Non-Financial Risks: Model Risk	9
2	Quantifying the Liquidity Premium on Corporate Bonds	14
2.1	Introduction	14
2.2	Review of Literature	17
2.2.1	Liquidity Proxies	20
2.2.2	Empirical Results	24
2.2.3	CDS-based Approach	24
2.2.4	Structural Model Approach	28
2.3	Description of the Data	37
2.4	Descriptive Analysis of Credit Spreads	43
2.4.1	Timeline Analysis	45
2.4.2	Descriptive Analysis of Market Data	48
2.5	Modelling Process	52
2.5.1	Modelling Methodology	53
2.5.2	Modelling the Bid-Ask Spread	54
2.5.3	Modelling the Credit Spread	55
2.5.4	Credit Spread of Perfectly Liquid Bonds	56
2.6	Numerical Results	57
2.6.1	Modelling Bid-Ask Spread	57
2.6.2	Modelling Credit Spread	62

2.6.3	Liquidity Premium Estimates	66
2.7	Additional Analyses	70
2.7.1	Alternative Liquidity Proxy Specification	71
2.7.2	Robustness of Credit Spread Measure	74
2.7.3	Investigating RBAS Properties	77
2.8	Bank of England’s Structural Model	80
2.8.1	Quick Model Overview	81
3	Quantitative Factor Investing in the UK Corporate Bond Market	84
3.1	Introduction	84
3.1.1	Review of Credit Specific Literature	87
3.2	Description of the Data	90
3.3	Methodology & Risk Factors	91
3.3.1	Risk Premia Factors	92
3.4	Single-Factor Portfolios	98
3.4.1	Non-Normality of Factor Returns	98
3.4.2	(Risk-adjusted) Measures of (Out-)performance	100
3.4.3	(Risk-adjusted) (Out-) Performance of Factor Portfolios	102
3.4.4	Turnover & Transaction Costs	107
3.4.4.1	Choice of Factor Tolerance	108
3.4.4.2	Annualised Performance Measures	110
3.4.5	Illiquidity Factor as Liquidity Premium	112
3.5	Multi-Factor Portfolios	114
3.5.1	Benefits of Natural Crossing and Reducing Trading Costs	123
4	Model Uncertainty & Parameter Risk in Stochastic Credit Models	124
4.1	Introduction	124
4.2	Review of Modelling Methodologies	127
4.2.1	Review of Credit Ratings as a Markov Chain	128
4.2.1.1	Finding a Generator for Public Migration Matrices	131
4.2.1.2	From Published Matrix to Valid Generator	132

4.2.1.3	Historical Data	135
4.2.2	Review of Default Probability and Credit Spreads	137
4.3	Ignoring Complexities in the Rating Process	143
4.3.1	Rating History	143
4.3.1.1	Evidence for Downward Momentum	146
4.3.2	Time Inhomogeneity	148
4.3.2.1	Evidence for Time-Inhomogeneity	151
4.3.3	Interaction between Rating Drift and the Business Cycle	155
4.3.4	Statistical Uncertainty	157
4.3.5	Stochastic Credit Model & Calibration	159
4.3.6	Calibrating the Market Price of Risk	160
4.3.6.1	Calibration of a CIR process	161
4.3.6.2	Likelihood Function	162
4.3.6.3	Initial Estimates using OLS	163
4.4	Empirical Results using Monte Carlo Simulation	164
4.4.1	Simulation of Rating Process	164
4.4.2	Empirical Results of the Simulated Portfolios	166
4.4.3	Sensitivity Analysis	171
5	Closing Remarks & Directions for Future Research	173
5.1	Quantifying the Liquidity Premium on Corporate Bonds	173
5.2	Quantitative Factor Investing in the UK Corporate Bond Market	177
5.3	Model Uncertainty & Parameter Risk in Stochastic Credit Models	181
A	Appendix to Chapter 2	184
B	Appendix to Chapter 3	186
C	Appendix to Chapter 4	188
	Bibliography	189

List of Tables

2.1	Joint distribution of successive price changes (conditional on no new information).	21
2.2	Average of daily standard deviations by variable and rating class. Note that Duration and Notional Amount are logged variables in the model.	58
3.1	Risk and Return measures (excess returns, gross of transaction costs) for all factor portfolios over the period October 2003 - July 2014. . .	104
3.2	Monthly turnover and excess returns (gross and net) over duration-matched Gilts for each of the factor portfolios.	108
3.3	Annualised volatility estimates using different scaling parameters; scaling under the assumption of independence, and the scaling factor under the assumption of an estimated AR(1) process.	112
3.4	Risk and return characteristics of Liquidity factors constructed on subsets of the bond universe data.	113
3.5	Pearson correlation between excess returns over duration-matched treasuries for factor portfolios and the market portfolio.	115
3.6	Pearson correlation between the outperformance over the market portfolio for all seven factor portfolios.	115
3.7	Optimised factor portfolio weights in percentages under each of the strategies.	119

3.8	Comparing two simple weighting schemes, equal weighted (W_E) and inversely proportional to variance (W_I), to the risk and return metric using the optimal portfolio weights (W_O). Each scenario is evaluated here using a difference risk or return metric; the portfolio weights W_O vary for each of the scenarios, whereas the weights for W_E and W_I are fixed by definition.	120
3.9	Portfolio weights for the optimal portfolio and the portfolio that meets the similarity criterion, but is least similar, for each of the strategies.	121
4.1	Migration matrix published by Moody's in their annual default study (Moody's Investor Services, 2011)	133
4.2	Comparison of methods of conversion from published annual migration matrix to generator and back.	135
4.3	One-year CTMC benchmark migration matrix estimated on all rating events from US-based issuers (1980 - 2002).	137
4.4	Possible migrations under the extended state space are marked with an 'X'.	145
4.5	Long term average migration matrix of 5-year transition probabilities published by Moody's in their annual default study (Moody's Investor Services, 2011).	146
4.6	Best estimates of the default probability for all ratings and states and the corresponding 99% confidence interval.	148
4.7	Persistence and default probability conditioned on the state of the economy across rating categories.	154
4.8	Best estimates of the default probability and the corresponding 99% confidence interval.	158
4.9	Quarterly transition matrix for three economic regimes, defined by GDP growth.	169
4.10	Value-at-Risk relative to the benchmark matrix with standard model parameters (Benchmark).	171

B.1	Long-term credit migration matrix, estimated by Service in their Annual Default Study of 2014.	186
C.1	Sector and geographic areas of firms included in the dataset (%). . . .	188

List of Figures

2.1	The Annual Benchmark Spread is based on the difference in annual yield between a bond and the corresponding benchmark bond. A benchmark bond is chosen as to minimise the difference in maturity (remaining life) between a bond and its reference.	43
2.2	Daily average Credit Spreads (and 5 th and 95 th quantile) over time for Financial/Non-Financials and the four investment grade rating categories.	44
2.3	Average Credit Spread by Rating (on a logged y-axis) over time, annotated with a non-exhaustive selection of events during the ‘credit crunch’. Tick marks indicate the start of the year.	45
2.4	Descriptive analysis of market data: number of issuers (top left), total notional amount outstanding (top right), average duration by rating (bottom left and the proportion of the bonds universe issued fewer than twelve months ago.)	48
2.5	Exploratory Analysis of a selection of available analytical values and their change over time (top), coupled with the way in which some bivariate relationships vary over time (centre, bottom).	50
2.6	Stylised representation of bond yields, illustrating the challenge to estimate yield curve C in order to extract liquidity premia. Yield curves: A- risk free (e.g. gilts); B- as A plus expected default losses; C- as B plus credit risk premium; D _{1,2} - as C plus liquidity premium and bid/ask spread; E _{1,2} - as D but higher bid/ask spread.	53
2.7	Beta parameters (β) for log duration, grouped by Financials and Non-Financials.	58

2.8	Beta parameters (β) for log Notional (left) and Seniority indicator (right).	59
2.9	Beta parameters for the Age Category indicator (left) and Capital Tier (LT2) indicator (right).	60
2.10	The weekly duration coefficient, β , for Non-Financial Issuers (left) and Financials (right) for all four rating categories.	61
2.11	Gamma coefficient for Non-Financials (left) and Senior (right) bond indicators.	62
2.12	Gamma coefficient for Non-Financials across rating categories.	63
2.13	Gamma coefficient for log Duration (NF) on the left and RBAS on the right.	64
2.14	Gamma coefficient for Coupon (left) and Collateralized indicator (right).	65
2.15	Variation of the coefficient of determination, R^2 , for the credit spreads model over time (weekly) and by rating category.	66
2.16	Decomposition of credit spread (left) for A-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).	67
2.17	Time-varying nature of four weekly quantiles of daily liquidity premium distributions, in basis points (left) and in proportion of spread (right).	68
2.18	Monthly estimates of median liquidity premia across rating category.	69
2.19	Predicted liquidity premia for constant bond over time (left), for various degrees of liquidity (right).	70
2.20	Shown on the same scale, the bid-ask spread for all IG ratings, both Financial and Non-Financial issuers, increased dramatically during the financial crisis.	72

2.21	Comparison of A-rated liquidity estimates in basis points (left) and % of spread (right) for the proposed modelling approach (black) and the alternative approach (grey).	74
2.22	Daily average values of three credit spreads measures for Non-Financials across four rating categories.	75
2.23	Daily average values of three credit spreads measures for Non-Financials across four rating categories.	76
2.24	Comparison of liquidity premia estimates across various credit spread measures, aggregated by year to aid interpretation.	77
2.25	Bid-ask spreads (left) and RBAS (right) for bonds issued by selected Financials. Ordering of bonds (by issuer and then by magnitude of BAS) on the left is preserved in the right-hand plot.	78
2.26	Bid-ask spreads (left) and RBAS (right) for bonds issued by selected Non-Financials. Ordering of bonds (by issuer and then by magnitude of BAS) on the left is preserved in the right hand plot.	79
2.27	Bid-Ask Spreads (left) and RBAS (right) for three individual bonds over time.	80
3.1	Boundary value M_t indicating the highest time-to-maturity included in the Quality factor on each day. Tick marks indicate 1 st January.	95
3.2	Distribution of returns for each of the factor portfolios and the market portfolio. The blue and red line correspond to the fitted densities of a skewed t - and Gaussian distribution respectively.	99
3.3	Total return for each of the factor portfolios and the market portfolio (left) and the total return in excess of duration-matched Gilts (right). Tick marks indicate the start of the year.	103
3.4	For various levels of factor tolerance λ , the turnover (left), gross excess return (middle) and net excess return (right), for each of the seven factor portfolios.	109
3.5	Optimisation of multi-factor approach for five scenarios describing an institution's risk/return objective.	118

3.6	Bootstrapped distribution of out-of-sample outperformance versus the market for optimal weights (strategy A1), equal weights and weights inversely proportional to volatility.	122
4.1	Sample of rating events (US) broken down over time and sector.	136
4.2	Real-world default probabilities (left) and corresponding credit spreads (right) against the term to maturity. The seven curves correspond to seven initial credit ratings; AAA to Ca-C.	138
4.3	Default probabilities (left) and corresponding credit spreads (right) against the term to maturity, introducing a market price of risk $\mu_{ij} - 1 = 0.4$ for all i and t . The seven curves correspond to seven initial credit ratings; AAA to Ca-C.	141
4.4	Risk Premium as a function of Term to Maturity, where the seven curves correspond to seven initial credit ratings; AAA to Ca-C.	142
4.5	Bootstrapped distribution of 1-year default probabilities for B (left) and B* (right) rated bonds.	149
4.6	Division of sample period by NBER recession indicator (left) and real GDP growth (right).	150
4.7	The up/downgrade ratio (left) and the number of defaults (right) vary over the sample period.	152
4.8	Estimates of default probabilities by rating fitted on annual subsets of the dataset.	153
4.9	Rating persistence across all rating categories for ten GDP growth buckets.	155
4.10	Difference in default probabilities between excited and non-excited counterparts for the CTMC fitted to the entire sample, and the recession/non-recession subsamples. The difference in PD (x-axis) is provided on a log-scale.	156
4.11	Bootstrapped distribution of 1-year default probabilities for BBB (left) and CCC (right) rated bonds.	158

4.12	VaR differences from benchmark matrix for simulation with standard parameters.	168
4.13	Relative distribution of rating membership at various years ahead, for the benchmark model and the economy dependent model	170
5.1	Schematic, theoretical decomposition of Credit Spreads, with a time-dependent element.	176
A.1	Decomposition of credit spread (left) for AAA-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).	184
A.2	Decomposition of credit spread (left) for AA-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).	185
A.3	Decomposition of credit spread (left) for BBB-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).	185
B.1	Number of bonds in each of the factor portfolios that disappear from the dataset altogether, but have not matured.	186
B.2	Number of bonds in each of the factor portfolios that mature in any given month.	187
B.3	Number of bonds in each of the factor portfolios fall outside the factor exposure criteria, using a tolerance level of $\lambda = 40\%$	187
B.4	Sample autocorrelations for the monthly log returns of each factor portfolio.	187

Chapter 1

Motivation to the PhD Thesis

This thesis is the end-product of my PhD research on the topic of liquidity, factor investing and model uncertainty, specific to the corporate bond market in the United Kingdom. All topics of research covered in this thesis have been, or are currently, ‘hot’ topics affecting the corporate bond market, whether purely from an insurer’s perspective, a regulatory perspective or a general investor’s perspective. In the section that follows, an introduction to the topics in this thesis and how it fits within the body of literature, are provided. Before discussing the motivation and background to the research presented in this PhD thesis, I feel it is necessary to mention how this thesis came about. This research was conducted under the umbrella of the Actuarial Research Centre (ARC) of the Institute and Faculty of Actuaries (IFoA) and is the result of one of its two inaugural projects. Financial assistance was provided by Partnership Assurance Ltd., with support from jointly the IFoA and ARC-appointed supervisors from Heriot-Watt University. The purpose of the PhD studentships under the ARC are to encourage collaboration between academia and industry to generate impact. Partnership is a specialist provider of (enhanced) retirement annuities with more than 5 billion pounds under management, looking after affairs for more than 120,000 policyholders. The need for in-house, longer term research into the liquidity and investment practices in the domestic bond market brought Partnership and the ARC together. As a result of the strong desire to generate insights for the benefit of Partnership and industrial partners, the topics covered in this thesis share a practical and data-centric philosophy.

1.1 Regulatory Environment of the Liquidity Premium in Solvency II

In their Consultation Paper 40, the Committee of European Insurance and Occupational Pensions (CEIOPS) discussed the discount rate used to calculate the present value of future cash flows of insurance obligations for the purpose of setting technical provisions. With the release of the consultation paper, concerns were raised because of the absence of a liquidity premium. Suggesting that the lack thereof would increase the cost of providing annuities in the UK market, news outlets reported that this cost would subsequently be passed on to newly written annuities.

Where assets are illiquid, investors demand an additional premium as a reward for the risk of incurring additional costs and for the uncertainty about costs in case the asset has to be sold at a future date. This additional premium leads to an increase in the implicit yield of the instrument. However, the liquidity premium is only one component of the total spread between the yield of a credit-risky asset and the liquid risk-free rate. This spread also includes a compensation for other risks such as expected credit losses, credit risk (unexpected credit risk/losses) and a ‘residual’ element, which could represent compensation for anything from taxes to conversion costs.

The cash flows of annuity contracts are more predictable than other types of insurance contracts and therefore easier to match using trade-able assets. Generally, gilts and corporate bonds are used to match annuity cash flows and these bonds will be on the books of an insurer with the intent to hold until maturity. The bondholder can mitigate against the risk attached to the part of the credit spread attributed to liquidity in a straightforward way; holding the bond until maturity. Then, the argument goes that the bondholder, in this case the annuity provider, can earn the value of the liquidity risk premium in their valuation.

Thus, to determine the part of the spread attributable to liquidity risk, the challenge that has to be faced is the accurate breakdown of this spread into its components. In this thesis, the focus lies with the research to establish, using data-driven

methods and practical for potential industry implementation, the size and dynamics of the compensation for the ‘illiquidity component’ of corporate bond spread. This thesis does not attempt pass judgement, nor does it attempt to quantitatively describe or stress test values of asset and liabilities under any of the proposed pieces of legislation that discusses a liquidity premium, adjustment factors or discount rates of the last five years. The remainder of this motivation briefly covers the timeline of the regulation and considers the way the Matching Adjustment addresses the issue of the “liquidity premium” before coming back to the core of the research to review briefly the types of analyses that are being used to disentangle the credit spread and how this thesis contributes to that discussion.

In their final advice in the consultation, CEIOPS concluded further work on the liquidity premium for technical provisions was necessary and a Task Force was created. The European Commission requested CEIOPS to run the fifth Quantitative Impact Study (QIS 5) for Solvency II, which was published in March 2011. Subsequent important milestones in the evolution of the liquidity premium include the Long Term Guarantees Impact Assessment, published in June 2013 by the European Insurance and Occupational Pensions Authority (EIOPA) and the Omnibus II directive that passed European Parliament in June 2014.

- **QIS 5** The liquidity premium is recognised and can be calculated using a straightforward formulaic approach: $LP_{assets} = \max(0, x \times (\text{Spread} - y))$. QIS 5 has taken liquidity premium estimates from three pieces of external (practitioner) research and concludes that a simple approach can be calibrated to reach a good consensus among estimates; parameters of $x = 50\%$ and $y = 40\text{bps}$ are suggested. It notes the aim of practitioners to use the liquidity premium as a tool to mitigate pro-cyclical behaviour of the solvency balance sheet during market turmoil. It also includes a ‘three bucket system’, where the categorization of the liabilities into these three buckets depends on specific product features and the second bucket with an application of 50% of the illiquidity premium is applied where liabilities do not fall under one of the other buckets.

- **Long Term Guarantees Impact Assessment** Two new concepts are explored in this report; a Matching Adjustment (MA), that brings relief to qualifying liabilities backed by qualifying assets and a Counter Cyclical Premium (CCP) which could be activated to increase the discount factor, at EIOPA's discretion in the event of a sudden fall in financial markets. The MA is an allowance that can be applied to the discount factor when sets of criteria are met for the insurance liabilities and the criteria on the matching of its cash flows to assets held. The report considers several implementations of the MA such as the 'classic' and 'extended' Matching Adjustment. Regarding the CCP, the report expresses concerns over the triggering mechanism and recommends it is replaced with a 'volatility rebalancer' that would operate automatically.
- **Omnibus II** The expected cash flows from the MA Assets must replicate each of the expected cash flows of the MA Liabilities in the same currency, they must be fixed, cannot be changed by the issuer or any third party and there must be no mismatch giving rise to material risks. The size of the MA is determined by the (corporate) bonds held as MA Assets and are based on the spread less cost of expected defaults and downgrades (referred to as the Fundamental Spread), with a minimum of 35% of average long term spreads. The Fundamental Spread is set by EIOPA on a monthly basis.

Whereas the research in this thesis is motivated by a regulatory need (to break down the credit spread in liquidity and non-liquidity components), at its core lie statistical methods that are used to isolate the liquidity premium on corporate bonds. There is a considerable body of research that this thesis contributes to and the aim of the proposed modelling approach is to devise a modelling strategy that is straightforward enough to implement, relies on readily available data only and models liquidity at the level of individual bonds.

The modelling attempt in this thesis is far from the first study to concern itself with the disentanglement of the credit spread of corporate bonds. Chapter 2 discusses previous modelling efforts in great detail, highlighting how this thesis aims to contribute to the existing body of research. Models of the firm, or structural mod-

els, describe the dynamics of pricing corporate debt and find their roots in Merton (1987). The original Merton model has seen many extensions over the last decades, most of which attempt to investigate particular aspects of a firm in greater detail; some of these models are discussed and these innovations can pertain to incorporating stochastic interest rates, allowing for complex capital structures or allowing for endogenous default boundaries. The model innovations however, are pedagogical in nature and aim to highlight how particular characteristics of the firm, or assumptions in the model, translate to model outcomes such as estimated probabilities of default. Empirical evidence continues to show these models are capable of capturing the probability of defaults rather accurately, but fail to model credit spreads, especially at the level of individual assets.

The Bank of England published an influential working paper (Churm and Panigirtzoglou, 2005) where a structural model (Leland and Toft, 1996) was calibrated to the UK corporate bond market and the liquidity premium computed as the residual of model spreads and observed market spreads, computed at aggregate market levels. Other modelling efforts include CDS-based methods, where the reasoning is that the CDS premium is compensation for *only* the credit component of the credit spread which would make for an easy computation of the liquidity premium; unfortunately empirical evidence demonstrates this is not as straightforward. By far the largest group of models are looking at the extent to which illiquidity is priced in corporate bonds through statistical models. Generally, the literature of statistical models use regression analyses to relate some proxy for liquidity to observed bond prices/spreads, of which the most thorough treatment can be found in Dick-Nielsen et al. (2012). These statistical models vary widely in the time period under study, the sample of bonds that are investigated and the data requirements that come with the model; most importantly, the liquidity proxy that is chosen has tremendous effects on model outcomes. This thesis is of a similar nature and therefore suffers from some of the same pitfalls, but attempts to address several others explicitly; the model most notably constructs a new liquidity proxy, uses only readily available bond data and explicitly allows for the frequent estimation of bond-level premia.

1.2 Rise of ETFs and Smart Beta Products

Exchange Traded Funds (ETFs) have continued to present both challenges and opportunities for the traditional investment management industry as well as institutional and retail investors. Perhaps fuelled by persisting low yield in fixed income markets, the ETF hype also saw a tremendous increase in both the number of fixed income focussed ETFs and inflows into those ETFs, with fixed income ETFs seeing the largest inflows in January 2016 when investors poured 13 billion US dollars in US-based fixed income ETFs.

Exchange Traded Funds which are commonly referred to as falling under the umbrella of ‘smart beta’ funds, continue to be at the centre of a debate of active versus passive management and continue to grow in size, with 950 products currently on offer and more than 475 billion dollars invested (December 2015 estimates by Morningstar (Johnson et al., 2016)).

The first index trackers surfaced in the early 1970s and tracked the performance of DJI 30 stocks, on which Malkiel (2007) commented in his book (first published in 1973) that *‘What we need is a no-load, minimum management-fee mutual fund that simply buys the hundreds of stocks making up the broad stock-market averages and does no trading from security to security in an attempt to catch the winners. Whenever below-average performance on the part of any mutual fund is noticed, fund spokesmen are quick to point out ‘You can’t buy the averages.’ It’s time the public could’*. Standard index trackers do exactly this, with very low expense ratios. Research by Morningstar (Johnson et al., 2016) suggests that strategic beta ETFs are on average more expensive than their market-cap weighted counterparts with a total expense ratio up to three times higher on average, but note that fees have been falling and are expected to continue falling as the market for smart products matures and competition intensifies.

One could argue that the ‘active’ element of strategic-beta ETFs justifies higher fees, but investors should remember that fees are a great predictor of future performance over the long term and that the costs of running smart indices are not necessarily significantly higher than the costs of maintaining a standard market

weighted fund. Therefore, it is reasonable to question whether the higher costs are justified and investors need to consider carefully whether their belief in the merits of the underlying strategy over long time horizon justifies the higher expense.

The roots of ‘smart beta’ products can be traced back all the way to early days of Modern Portfolio Theory (Markowitz, 1952), where holding negatively correlated assets reduced a portfolio’s risk. In the 1960s, this led to the development of the Capital Asset Pricing Model (CAPM), where the concept of ‘beta’ first made an appearance. Representing the systemic risk and return of a fully diversified market portfolio, this beta is generally captured by passively managing a portfolio against a market-capitalisation weighted (market cap) benchmark. Later, Jensen (1968) introduced ‘alpha’, which represents the outperformance. If beta is delivered through passive management and alpha through active management, where does ‘smart beta’ come in?

Weights according to market capitalization might not necessarily describe a ‘fully diversified’ investment as suggested in a CAPM world. Currently, Apple and Microsoft account for more than 20% of the NASDAQ 100, even though the index consists of 100 stocks, these can be considered individual weights that represent an overweighting towards any single asset. Consequently, a vast array of smart beta strategies have been built, typically involving analytics and screening to build passive portfolios of assets weighted by almost any criteria other than market cap. Rather than an attempt to create a better way of building a fully diversified investment by using alternative weights, these funds might strategically expose the resulting portfolio to underlying factors/constructs. This strategic exposure can be interesting (for generating traditional alpha) and investors in ETFs can easily expose their portfolio to risks in order to achieve or enhance, for example, the risk and reward trade-off they are after. Since alpha is considered a measure of outperformance, it is common practice to judge investments and portfolio managers by their alpha statistics as an indication of ‘true’ skill, but measurement of alpha is far from straightforward. Most importantly, one needs to distinguish between alpha and beta, a line which can be very fuzzy indeed. This performance measurement has

been the subject of many discussions, where Kidd (2014) very eloquently describes how certain investment strategies, capitalizing on systematic return drivers ‘occupy the space between traditional beta and alpha’.

This thesis aims to capture the effect of some well-documented factor risk premia (systematic return drivers) from (international) equity markets, applied to the corporate bond market. These phenomena are well-established in academic literature and are the building blocks or philosophy of many traded alternative/strategic/smart beta ETFs and include risk drivers such as Value, Small Caps and Low Volatility. Not only are comprehensive reviews of the risk and return characteristics of several factor premia in the corporate bond market sparsely researched, the thesis also contributes to the existing body of research by quantifying the impact of corporate bonds trading in an illiquid market, bearing non-negligible transaction costs. To further the treatment of quantitative factor investing in corporate bonds, the thesis challenges the typical fixed holding period (6-12 months) under which strategies are typically evaluated, with allocated opportunities for rebalancing. Rather, a flexible rebalancing model is constructed that aims to maximise returns net of transaction fees as a result of diluted factor exposure and decreased portfolio turnover.

The methodology for arriving at factor portfolios, which is the ‘secret sauce’ in *true* smart beta funds, is extremely simple in definition, yet not necessarily out of touch with reality. Using some results from the statistical investigation earlier in the thesis, we use the attractive properties of the newly created liquidity proxy to build illiquidity factors that will not only give us insight into the performance of underlying constructs, but might also be used to estimate (market-wide) liquidity premia.

Special attention is given to an illiquidity factor definition as an alternative way of deriving the liquidity premia of subsets of the investment grade market. Based on simple (‘dumb’) specifications of the factor portfolio, a nearly model-free approach is used to see whether estimates of liquidity premia arise by observing the risk and return characteristics of the portfolio.

1.3 Financial Risk of Non-Financial Risks: Model Risk

Mathematics has been the cornerstone for innovation in the financial sector, from the Nobel prize winning Black-Scholes formulae (Black and Scholes, 1973) to the bundling and trading of assets in the early 2000s. Regardless of the root causes and catalysts of the financial crisis, the aftermath of the financial crisis saw repeated calls for the urgent need to gain better insights into the risks of mathematical modelling. Articles in popular news outlets (Salmon, 2009) reporting that Li's (1999) Gaussian copula is 'The Formula that Killed Wall Street', fail to distinguish between the merit of a mathematical construct, the assumptions in applying a mathematical theory and the misjudgement or disregard of the limitations presented by mathematical model. The aftermath of the crisis has most definitely led to increased scrutiny of mathematical modelling, which can only be encouraged, but in itself passes no judgement on the use of mathematics in finance.

Yet, both industry practice and regulatory oversight frequently fail to acknowledge adequately the risk that models themselves carry, despite this issue being highlighted for a considerable time (examples include Hendricks, 1996 and Berkowitz and O'Brien, 2002). After all, the existing literature is relatively sparse on quantifying model risk and generally only looks at back-testing as a means of assessing model risk.

Model risk and uncertainty should be an integral part of the risk management function and whereas quantifying model risk can be difficult, the concept of model risk is intuitive to define. Essentially, all financial models are wrong in the sense that they represent a simplification to a real-world phenomenon. However, there are many potential models to describe any particular phenomenon of financial markets including option pricing, hedging strategies or portfolio allocations; the simplification depends on the chosen model. To describe a real-world phenomenon, one firstly has to decide which model to use from a large pool of potential and seemingly valid candidate models. After choosing a specific model, one needs to decide on model pa-

rameters, which in turn can either be entirely subjective or can be prone to indirect subjectivity as one needs to choose calibration methods and calibration/validation data. This leaves the modeller exposed to both model and parameter risk. Chapter 4 aims to study *some* of the model and parameter uncertainty that comes with projecting forward credit-risky instruments using a stochastic credit model.

Model risk and model uncertainty, or parameter risk and parameter uncertainty are commonly used interchangeably (as are the more general concepts of risk and uncertainty). In this thesis these terms are mostly used interchangeably in the empirical work, but it is worth making a distinction between the two concepts, based on the seminal work published by Knight (1921). His work distinguishes between two situations;

- The probabilities of every possible outcome are known
- The probabilities of every possible outcome are unknown

The latter situation is referred to by Knight (1921) as uncertainty whereas the situation in which a probabilistic description of outcomes can be made, is referred to as risk. Naturally, risk is a special case of uncertainty and a far better situation to be in than uncertainty, as risk allows for risk management, that is, taking advantage of the known information about probabilities to act in a way that minimizes some function concerning adverse effects.

Extending the concepts of risk and uncertainty to financial models, a situation where a modeller has to choose a model (P) from a set of possible models (\mathcal{P}) would be referred to as *model uncertainty*. If every model $P \in \mathcal{P}$ can be described by the parameter set θ from a space of parameters (Θ), one speaks about *parameter uncertainty*. However, if a probability measure R is available on \mathcal{P} on the set of possible models (or parameter space Θ) that indicates the model as the correct model, one deals with the *model risk* (or parameter risk).

Whereas mathematical models of the physical world date back hundreds or thousands of years in physics, mathematical models and stochastic models in particular, are a far more recent trend in finance, often credited to Bachelier (1900). Comparing

the modelling of financial markets and the modelling of classical mechanics, one can easily observe the tremendous complexity of the financial markets given the many forces that are at play, and the influence of those forces are non-negligible and ever-changing. One could argue that given the endless possibilities to simplify financial processes, many different models are competing with each other at all times, leading to (true) model uncertainty. Parameter uncertainty is more difficult to pin down as in most models, it is likely that some probabilistic measure exists around the true parameters, which would lead to model uncertainty and parameter risk according to the classification by Knight (1921); hence the title of this Chapter.

Whereas the focus in Chapter 4 lies with the model uncertainty and parameter risk in stochastic credit models and specifically the risks associated with the use of readily-available migration matrices, the modelling of liquidity premia in Chapter 2 also touches upon the topic of model risk heavily. Chapter 2 reviews many implementations of statistical models and structural models that have been used to extract liquidity premia. In particular it stresses not only the model uncertainty of trying to estimate liquidity premia, but also the parameter risk of structural models; these models have many, rather complex parameters, which makes for a subjective calibration process.

With the importance of model and parameter risk highlighted and the concepts defined, the question arises how to study the model and/or parameter risk of a modelling challenge. Danielsson et al. (2015) argue that back-testing is a relatively poor method to capture model risk as it is highly dependent on assumptions about the statistical distribution of financial variables, for which, almost by definition, insufficient information exists. In the same work, they argue that looking at the disagreement or discrepancies amongst candidate models is a far superior method of assessing model risk. They find that market risk models exhibit model risk that becomes particularly apparent during periods of financial turmoil. They argue that this is particularly concerning as risk models are designed to manage the risks of crisis. However, not only is market crash data rare and models would struggle to capture these infrequent events, the risk management function is also likely to place

an emphasis on the management of day-to-day risks. Danielsson et al. (2015) note however that of particular concern they find that financial turmoil is considered exogenous to the model.

The model risk assessment in this thesis is concerned with stochastic credit models and takes a similar approach to Danielsson et al. (2015) in that the ultimate outcomes of a typical simulation exercise are compared against the outcomes using several competing models. The competing models are investigated for the risk associated with one aspect of credit models only; the specification of migration matrices that form an integral part of the model.

When used in simulation exercises, the migration matrices are vital to the stochastic evolution of an obligor's credit quality and the matrices are seen as an exogenous input, generally taken directly from rating agencies' publications. The matrix used is generally the non-sector specific, long term (30 year) average matrix, as computed and provided by the rating agency. There is a disregard for the modelling that is used to arrive at rating matrices or the assumptions underlying the matrix as a Markov Chain to project forward credit ratings.

Using raw data, that is, rating event data, this thesis attempts to quantify the model risks embedded in typical stochastic credit models and refers to these risks as *risks of ignorance*. Investigating the risk of ignoring two non-Markovian properties of the rating process that are well-established (time homogeneity and rating momentum) and additionally investigating the risk of ignoring statistical uncertainty in the best estimates of the migration matrices, model risk is quantified by comparing risk measures of simulated credit-risky bond portfolios across competing models.

The research in this thesis may be relevant to any application of stochastic credit models and may be of particular interest to projections of spreads over long time horizons, such as the simulations from an Economic Scenario Generator (ESG). For companies relying on solutions provided by external ESGs, the analysis might provide some guidance to interpreting the uncertainty in the produced estimates or might serve to start meaningful discussions regarding the model risk with their provider. For those building their own stochastic credit models, the analysis illus-

trates the potential added value of incorporating the modelling of the migration matrix from first principles rather than relying on migration matrices that are readily available. Not only can the model be better tailored to its specific purpose, it allows for a comprehensive review of the some of the embedded risks in the model.

Chapter 2

Quantifying the Liquidity Premium on Corporate Bonds

2.1 Introduction

In 2009 the Committee of European Insurance and Occupational Pensions (CEIOPS) set up a special task force to investigate the Illiquidity Premium. In March 2010 the task force published its findings and argued for the introduction of an illiquidity premium. The argument made by the insurance industry is simple in principle. The cash flows associated with annuity contracts are predictable and easy to match with traded assets; gilts and corporate bonds are used to match liability cash flows and will be held until maturity. The yield on a corporate bond is higher than the yield on a ‘risk-free’ gilt or swap, with the additional yield referred to as the credit spread. This credit spread compensates investors for excess risk the holder bears; for example, the issuer of the bond may default, the issuer of the bond may get downgraded and the bond may be difficult to convert to cash, i.e. it is illiquid. Whereas the default related risk is always present and cannot be fully diversified away, the difficulty to turn a bond into cash is easily avoided by holding the bond until maturity. The argument goes that the bond holder can capture the value of that liquidity risk premium in their valuation. The difficulty lies in calculating or even defining the liquidity of insurance liabilities, as it cannot be easily compared to the liquidity of assets. For some sceptics this fundamental issue is where the

insurer's argument breaks down. On the other hand, CEIOPS defines the liquidity of liabilities as the degree to which a liability's cash flows are predictable, since predictable cash flows can be replicated using illiquid assets. This implies that the definition of 'illiquid' liabilities excludes liabilities for which cash flows are not as predictable (because the policyholder can surrender their policy before the redemption date for example). Determining the eligibility of liabilities based the proposed definition of 'predictability' is problematic since the relatively unpredictable cash flows, can nonetheless be matched with a combination of bonds and derivatives that can be held to maturity. In its first report in 2010, the CEIOPS task force states that 'to determine the part of the spread attributable to liquidity risk, the challenge that has to be faced is the accurate breakdown of this spread into its components' (CEIOPS, 2010). In the same report it offered a simple formula that could be used as a proxy for the Liquidity Premium, based on a fixed percentage of the spread. The calibration of this one parameter would determine the liquidity premium and issues were raised immediately to the potential subjectivity and the frequency of calibration.

The technical issue of quantifying the Liquidity Premium on corporate bonds is the focus of this Chapter. In particular, this Chapter focuses on a review of popular (mathematical) methods of corporate spread decompositions and assesses their results and practical use. Previous work regarding spread decompositions had four major issues making industry adoption problematic:

- outcomes relied heavily on chosen sample of bonds, time period or require the calibration of many parameters
- modelling techniques required input data that is simply unavailable for a large part of the relevant bond universe
- estimates of liquidity premium are aggregated and do not account for differences in for example rating, sector, duration or seniority and lastly
- all estimates are at low frequencies (monthly or quarterly)

Academic literature has studied the effect of illiquidity on corporate bond prices extensively over the past three decades, from both theoretical and empirical perspectives. Amihud et al. (2006) discuss a series of asset pricing models in which frictional costs lead to higher expected returns, compensating investors for investing in illiquid assets. The work in Amihud et al. (2006) is a special case of Amihud and Mendelson (1986) where investors have exogenous time horizons and assets bear illiquidity due to exogenous trading costs. Amihud et al. (2006) and Acerbi and Scandolo (2008) also discuss the heterogeneity of corporate bond investors with respect to expected holding times and how these different groups lead to a market equilibrium in which investors with short expected holding periods hold very liquid assets and investors with the longest expected holding periods hold illiquid assets.

A related concept in asset pricing theory is that of the marginal investor that ultimately determines an asset's price (Sharpe, 1964; Cochrane, 2005). With respect to the corporate bond market, this raises the question whether there are sufficient hold-to-maturity investors to take up the entire supply of corporate bonds; if sufficient long-term investors are vested in the market the yield spreads would only reflect credit factors and liquidity premia would be very small.

Empirical literature investigates whether illiquidity is priced by relating bond prices to various proxies for liquidity using a reduced form modelling approach, but also aims to quantify the liquidity premium as part of the credit spread. In addition to using reduced form models to quantify the liquidity premium, structural models of default (for example Merton, 1974 and Leland and Toft, 1996) have been used, as have direct computation methods (for example Breger and Stovel, 2004 and Koziol and Sauerbier, 2007). Section 2.2 discusses the empirical literature in more detail.

With the focus solely on the quantitative measurement of liquidity premia in the corporate bond market, this Chapter leaves the regulatory aspects of market consistent valuation, matching assets, discount factors and pro-cyclicality far behind. In this Chapter, the liquidity premium is defined as the difference in yield to maturity of a bond relative to the yield on a hypothetical perfectly liquid bond with otherwise identical characteristics. To an investor who is prepared to buy and hold to maturity,

the liquidity premium represents the expected reward, per annum, in return for sacrificing the option to sell a bond before maturity. Any investor who plans to or might need to sell before maturity will, on average, earn a lower premium than the estimated liquidity premia.

A new methodology for estimating liquidity premia on corporate bonds is formulated, addressing some of the pitfalls of other modelling approaches. Using quoted bid-ask prices and a comprehensive dataset (Oct 2003 - Jul 2014) of end-of-day bond characteristics and statistics (GBP investment-grade), a liquidity measure, uncorrelated with bid-ask spreads and bond characteristics, is derived. This new liquidity measure is used to extract liquidity premia, but the liquidity score on individual bonds can be a useful tool in portfolio management in itself.

In addition to deriving a new liquidity measure, the research in this Chapter estimates liquidity premia on a more granular level than existing literature. Daily cross-sectional regression analyses allow estimation of liquidity premia at the individual bond level, on a daily basis. The research in this Chapter is novel in the sense that the same methodology for estimating liquidity premia is applied over a relatively long period of time (11 years), capturing both the benign economic climate prior to the financial crisis, the financial crisis, and more recent years. The modelling approach of daily cross-sectional regressions also allows for the evolution of model parameters to be studied. In particular, the time-varying nature of liquidity premia, both in basis points and in proportion of total credit spread, is clearly visible.

2.2 Review of Literature

Liquidity is regarded a desirable property of both financial markets and asset classes in aggregate and individual securities. Before delving into attempts at defining liquidity, it is important to distinguish between two ‘types’ of liquidity. Liquidity in financials markets can generally refer to either liquidity in funding, or liquidity in trading. Following Brunnermeier and Pedersen (2009), these two aspects of liquidity can be defined as follows;

- Funding liquidity is the ease with which market participants can obtain funding
- Market liquidity is the ease with which an asset can be traded.

Even though this Chapter is primarily concerned with so-called market liquidity rather than funding liquidity, the two cannot be seen in total isolation, especially during periods of (extreme) market distress when funding and market liquidity can lead to a spiral of illiquidity (Brunnermeier and Pedersen, 2009). Brunnermeier and Pedersen (2009) argue that their mutually reinforcing effect during crisis periods causes well-known liquidity phenomena such as ‘flight-to-quality’. When traders (speculators, hedge funds, investment banks, all marked-to-market) buy a security, they are required to use some of their own capital (difference between security’s price and its collateral value) to finance the trade. Similarly, short selling requires a margin on all positions. Traders are less willing to put on trades if funding is sparse, especially in capital-intensive securities, which has direct consequences for liquidity across the entire market.

Market liquidity is difficult to define, especially since different markets participants (for example, day traders and pension funds) and different sides of the market (buy or sell), might have different requirements of a liquid market. Various definitions are offered by recent literature and this section aims to extract all aspects of liquidity that might need to be considered when evaluating the liquidity of a market or individual security.

In its Stability Report in April 2007, the Bank of England (2007) argues that liquidity risk, defined as ‘the harmful consequences of illiquidity’, is present when ‘one cannot easily offset or eliminate a position without significantly affecting the market price’. Linking liquidity directly to adversarial movements in market price is likely to describe one aspect of liquidity, as the existing body of empirical research also often considers the price impact of trades when evaluating liquidity. However, on its own this definition is rather narrow in scope and many other aspects of liquidity are likely to come into play.

Taking a step back, the theory of market micro-structure studies market liquidity in the first instance. Market micro-structure theory is defined by O’Hara (1995)

as ‘the study of the process and outcomes of exchanging assets under a specific set of rules’. While much of economics originates from the mechanics of trading, micro-structure theory focuses on how specific trading mechanisms affect the price formation process. Kyle (1985) identifies three characteristics of a market that describe its liquidity:

- *Tightness*: size of the spread between bid-ask prices
- *Depth*: maximum trade size / volume that does not affect current prices, reciprocal of equilibrium price to trade volume
- *Resilience*: speed with which the price impact of trades disappear

Even though these three dimensions of market liquidity are unlikely to describe *all* aspects of liquidity, they are well-established. These dimensions form the basis of several factors that Amihud et al. (2006) identify as the main elements affecting the micro-structure of a market and the market liquidity of its traded assets: exogenous transaction costs, private information, inventory risk and search friction.

Exogenous transaction costs are costs incurred by both/either the buyer and/or seller of a security every time the security is traded. These costs can include transaction taxes, order processing costs and broker fees. Dealers will adjust their quote spread to protect themselves, on average, against counter-parties with superior knowledge from which a trading loss will occur. Inventory risk refers to the wider spread dealers will quote for holding an inventory in a security that deviates from their desired inventory and search friction describes how an investor might be making concessions on price if he cannot find a counterparty for his trade; specifically, it refers to the opportunity cost of an investor between an immediate transaction (price concession) and waiting for willing counter parties in the market.

The resulting bid-ask spreads are of particular interest, as they are central to the understanding and modelling of liquidity premia in this research. A different branch of the micro-structure literature, concerned with the decomposition of the bid-ask spread, considers all of the aforementioned components, such as exogenous transaction costs, private information, inventory risk and search friction. Early work

on decomposing the bid-ask spread focussed on explaining quoted spreads cross-sectionally using market variables such as trading volume and security risk (Benston and Hagerman, 1974 and Demsetz, 1968). Other work decomposes the spread into adverse information dealer profit components (Glosten and Harris, 1988 and Stoll, 1989), where the dealer profit represents compensation for inventory holding and order processing costs. Estimates by Stoll (1989) and Madhavan and Smidt (1991) indicate that inventory costs are relatively small for liquid asset classes but increase substantially for illiquid assets and, hence, the remainder of the spread is mostly determined by order processing costs and adverse information costs. Copeland and Galai (1983) observe that order costs, the clerical costs of carrying out the transaction, which includes the cost of the market makers' time, are fixed irrespective of trade size. Therefore, the average cost of order processing per unit decreases with trade size. The effect of informed counter-parties has been studied at length (see for example Glosten and Milgrom, 1985) and models describe how dealers / market makers demand compensation for losses incurred from counter-parties with superior knowledge. Specifically, Lin et al. (1995) empirically verify a model developed by Easley and O'Hara (1987) in which well-informed traders prefer to trade larger amounts.

With the above in mind, especially considering the 'search friction' identified by Amihud et al. (2006), another aspect of liquidity can be defined as 'immediacy'; the ability to execute a trade contemporaneously. Of course, immediacy is highly dependent on the trade size; at any given point in time, the ability to execute a trade with immediacy will differ substantially across small and large quantities. This implies that liquidity does not only operate at the level of (international) financial markets, varying by asset class or individual asset, but operates at a much more granular level; individual trades.

2.2.1 Liquidity Proxies

Empirical work involving liquidity and asset pricing relies on variables that are (indirectly) inherited from the micro-structure literature and are associated with

different levels of market liquidity. This section attempts to give an overview of ‘liquidity proxies’ that are extensively used in the empirical literature; by no means is the overview exhaustive in the proxies it lists.

The most intuitive liquidity proxy, and central in this Chapter, is the Bid-Ask spread, as it is regarded an ‘aggregate’ measure of liquidity (Hasbrouck and Seppi, 2001). The availability of bid-ask spread data is highly dependent on the choice of asset class under study, time period and granularity. An indirect, implicit approximation of the effective Bid-Ask spread introduced by Roll (1984), continues to be a widely used proxy in recent empirical work (Dick-Nielsen et al., 2012 and Bao et al., 2011). The Roll-measure constructs implicit Bid-Ask spreads based on prices alone. Recognising that negative serial dependence exists between observed price changes when dealing with a market maker, as first pointed out by Niederhoffer and Osborne (1966), Roll (1984) computes the spread as twice the negative covariance between subsequent price changes. To see why this negative dependence exists, Roll (1984) considers the following. Assuming, for simplicity, that all transactions are with a market maker with constant spread and no new information about the security arrives, successive transactions are equally likely to be a purchase or sale by the market maker since trades arrive randomly at both sides, exogenously. Therefore, the joint probability of successive price changes ($\Delta p_t = p_t - p_{t-1}$) depends on whether the last transactions was at the bid- or ask side. As the transaction at time $t - 1$ is equally likely to be at the bid- or ask side, the joint distribution of successive changes in price can be written using Table 2.1.

		Δp_t		
		-s	0	+s
Δp_{t-1}	-s	0	.125	.125
	0	.125	.25	.125
	+s	.125	.125	0

Table 2.1: Joint distribution of successive price changes (conditional on no new information).

Given that expected values of both Δp_t and Δp_{t-1} are zero and can be ignored,

the covariance reduces to;

$$\text{Cov}(\Delta p_t, \Delta p_{t+1}) = \frac{1}{8}(-s^2 - s^2) = \frac{-s^2}{4}$$

In addition to the Roll measure, other measures of (implicit and effective) Bid-Ask spread have been used. Other (implicit) spread measures include Holden’s (2009) effective spread measure based on observed price clustering and the “LOT” effective spread measure based on the assumption of informed trading on non-zero-return days and the absence of informed trading on zero-return days developed by Lesmond et al. (1999). Another popular measure of transaction costs inferred from transaction data is the Unique Roundtrip Cost (URC) measure, which matches up trades of similar volume within a short time period, assuming the trades occurred at different sides of the bid-ask. Contrary to literature that uses transaction data (Goldstein et al., 2007), primarily in the equity markets, the side of a particular trade as well as the type of agent executing the trade is unknown for Dick-Nielsen et al. (2012), using the TRACE (Trade Reporting and Compliance Engine) database for US corporate bonds.

A second class of liquidity proxies directly relates to market depth, one of market liquidity dimensions in Kyle (1985), and it assesses the price impact of trades. The idea is that a liquid asset should be able to trade in substantial quantities without moving price and, conversely, price movements will reflect the depth of a securities’ market. By far the most commonly used liquidity proxy in this class is the Amihud-measure (Amihud, 2002), which captures the ‘daily price response associated with one currency unit of trading volume’, serving as a rough measure of price impact. The defined ILLIQ measure (Amihud, 2002) dened as the average ratio of the daily absolute return to the (dollar) trading volume on that day; $\frac{|R_{iyd}|}{\text{VOLD}_{iyd}}$. Here, R_{iyd} is the return on stock i on day d of year y and VOLD_{iyd} is the respective daily volume in dollars. Defining ILLIQ annually;

$$\text{ILLIQ}_{iy} = \frac{1}{D_{iy}} \sum_{d \in D_{iy}} \frac{|R_{iyd}|}{\text{VOLD}_{iyd}}$$

where D_{iy} is the set of days for which data is available for stock i in year y .

The use of the Amihud-measure as a proxy for liquidity (price impact) is widespread in academic literature; from Dick-Nielsen et al. (2012) who use the measure in various regression analyses on the US corporate bonds market to Hasbrouck (2009) who links trading costs and returns in US equity markets. Its popularity and widespread use also caused the Amihud-measure to be subjected to further study, criticism and refinement. Theoretical work by Brennan et al. (2012) questioning the symmetric micro-structure framework suggested by Kyle (1985) finds that equilibrium rates of return are sensitive to changes between seller-initiated trades and returns, but not sensitive to buyer-initiated trades. Whereas the Amihud-measure treats positive and negative returns the same, Brennan et al. (2013) decomposes the traditional Amihud-measure into components that correspond to up-days and down-days, pointing towards research by, for instance, Brunnermeier and Pedersen (2009) arguing liquidity in a down market may be different from liquidity in an up market. They find that for US equity markets, the down-day component of the Amihud-measure is associated with a return premium whereas the up-day component is not significantly priced.

A third class of liquidity proxies can be referred to as trading intensity variables, which frequently covers both measures based on turnover and zero-trading-days. Turnover is intuitively defined as the ratio of the trading volume in a given period and the amount outstanding. The inverse of the turnover measure can be interpreted as the average holding period. Alternatively, zero-trading-days is a measure of trading intensity as it simply calculates the percentage of trading days with no trades in a given time period. In addition to a strict asset specific measure, Dick-Nielsen et al. (2012), analysing corporate bond liquidity, develop a firm specific zero-trading-days measure; the number of days in a given time period where none of the bonds issued by a particular firm trade. At any time, this measure tries to capture the fact issuers will have bonds of varying maturities outstanding and a shorter waiting time between trades within a firm indicates there is relatively frequent new information about the firm. In addition to liquidity proxies on their own, proxies can be aggregated into one

index or one component for their analysis. Kerry (2008) builds an index by averaging nine different proxies for liquidity including six micro-structure variables, evenly split between various bid-ask spread approximations and price impact (all return-to-volume). Dick-Nielsen et al. (2012) provide a comprehensive review of many liquidity proxies, all of which they subject to Principal Component Analyses to both assess communality between individual proxies and create ‘new’ aggregate liquidity measures. They conclude that the Amihud-measure and the Unique Roundtrip Cost measure are most consistent and statistically significant in explaining bond spreads, results that hold across bonds of different quality and market conditions (pre sub-prime crisis and during sub-prime crisis).

2.2.2 Empirical Results

The empirical estimation of liquidity premium generally considers one of three (categories of) methodologies for estimating this premium. The first methodology is a model-free approach where market prices for Credit Default Swaps are used, the second approach relies on the vast literature of the ‘value of the firm’ and the pricing of credit risky debt using option pricing theory and the last approach is referred to as ‘reduced-form’ and consists of empirical, statistical models that often rely on liquidity proxies in their estimation. The next sections will consider the CDS approach and structural models carefully by setting out the model rationale and methodology of literature in the category, evaluate empirical results and discuss the merits and imperfections of each presented approach.

2.2.3 CDS-based Approach

Credit derivatives cover a broad class of securitized derivatives whereby the credit risk of the underlying loan is transferred to an entity other than the lender (Sengupta, 2005). Credit derivatives can be divided into two groups; unfunded credit derivatives where two counter-parties sign a bilateral contract, and funded derivatives where the seller of protection puts up initial capital to settle possible credit events. Examples of unfunded credit derivatives include Total Return Swaps (TRS)

and Credit Default Swaps (CDS), the latter accounting for almost forty per cent of the total credit derivative market. Examples of funded credit derivatives include (synthetic) Collateralized Debt Obligations (CDOs). The market for CDS was non-existent before the 1990s and the creation of CDS as they are known today is often credited to JP Morgan & Co in 1994 (Stulz, 2009), selling on the credit risk from the line of credit it had extended to Exxon. The market growth of CDS since its inception is extraordinary, peaking at the end of 2007 when the notional value of the CDS market reached 62 trillion USD (International Swaps and Derivatives Association, 2010). Even when one acknowledges that the cash flow generated by the market is only a fraction of its notional value (historically only 0.2% of investment grade companies default in a year (Moody's Investor Services, 2011)), it has quickly grown to become an important financial market.

The structure of a single-name CDS contract is fairly straightforward; two parties are involved in the contract, the protection buyer who is looking to insure against the possibility of default on a particular bond and the protection seller, who is willing to bear the risk. The company that issued the bond is referred to as the reference entity, the bond itself being the reference issue. In case of a credit event (default, failure to pay or other 'trigger') the protection seller agrees to buy the reference issue at face value and in return receives a default swap premium, a periodic (quarterly) fee. The contract simply expires at maturity date in case no credit event happens during its lifetime and in case there is a credit event, the protection seller buys the reference issue at face value and the periodic payments are discontinued.

The CDS premium, the periodic fee at which the protection seller is willing to take on the risk of a credit event on the reference issue, is central to the Negative CDS Basis approach. Using arbitrage, Duffie (1999) shows that the spread of a corporate floating rate note (FRN) over a default free FRN should equal the CDS premium. Even though this is an approximation when applied to ordinary fixed coupon bonds, it is generally accepted. In reality, the difference between the CDS premium and the spread on the bond can be observed to be non-zero and negative, implying that other factors contribute to the entirety of the bond's spread. Longstaff and Schwartz

(1995), in their direct approach, interpret this negative basis as the difference in yield between an illiquid corporate bond (synthetically free of expected defaults and credit risk) and the yield on a liquid credit risk free bond. The residual yield is then interpreted as a direct quantification of the discounted yield associated with liquidity. This model free approach allows for the, in principal, easy computation of the liquidity premium using a simple equation:

$$\text{Liquidity Premium} = -\text{CDS basis} = \text{Corporate Bond Spread} - \text{CDS premium}$$

The extreme simplicity of the method comes with both a set of assumptions and data / computation restrictions. The key methodological assumption, based on Duffie (1999), is that the CDS premium is only compensation for bearing the credit risk of bond. In reality CDS contracts bear other risks, one of which is counterparty (credit) risk. One can argue that during periods of calm or benign financial markets, counterparty risk can be neglected, but in the aftermath of the 2008 global crisis, counterparty risk rapidly became a priority and concern. The extreme jumps in CDS spreads for troubled financials during the second half of 2008 seem to suggest that other factors, on top of mere credit risk, might contribute to the spread. The other, implicit, methodological assumption is that the negative basis is compensation for liquidity, whereas in fact, Longstaff and Schwartz (1995) are careful to point out that this is short-sighted.

Using a unique proprietary dataset with quotes and trades from fourteen CDS dealers selling protection on the same set of reference firms, Arora et al. (2012) investigate how counterparty credit risk affects CDS pricing. They find that despite the significant relation between dealer credit risk and cost of credit protection, the effect on CDS premia is small. Specifically, they estimate that a 645 basis point increase in dealer's credit spread translates into only a one basis point decline in selling the credit protection. Theoretical work exploring the magnitude of counterparty credit risk on CDS pricing generally estimates the price effect to be in the range of 7 basis points (Kraft and Steffensen, 2007) to 20 basis points (Hull and

White, 2000), implying an effect many times the empirical estimates. However, it is crucial to recognize that the theoretical literature focuses on CDS contracts in which the liabilities are not collateralized. Standard market practice during the sample period of the studies was full collateralization by both parties to the contract. Full collateralization would seem to imply that counterparty credit risk is not priced in CDS contracts. Reality is however, as became clear after the Lehmann bankruptcy, that firms putting down collateral in excess of their liabilities, often required by large Wall Street dealers (Arora et al., 2012), are at risk of becoming unsecured creditors of a defaulting counterparty. Arora et al. (2012) ultimately conclude that counterparty risk is most definitely priced in CDS contracts, but estimation cannot be seen without the context of collateralization, industry sector and chosen sample period.

Another factor that is likely to contribute to the CDS spread, beyond pure credit risk, is the illiquidity of CDS contracts themselves. Empirical work on the liquidity of credit derivatives is limited; Tang and Yan (2010) use regression analyses and capture the impact of expected liquidity and liquidity risk in CDS spread, Chen et al. (2005) use the term structure of CDS spreads to find both an expected liquidity premium and liquidity risk premium are earned, and Bongaerts et al. (2011) infer liquidity risk premia on CDS prices from expected excess returns. In other derivative markets, examples include Deuskar et al. (2011) who conclude that interest rate options with low levels of liquidity trade at higher prices than liquid equivalents and Cetin et al. (2006) who incorporate illiquidity into a standard Black-Scholes framework. Bongaerts et al. (2011) apply a theoretical asset pricing model that incorporates liquidity and allows for short selling to the credit default swap market. Estimating a non-linear asset pricing model by applying Generalized Method of Moments to quotes over a sample period between 2004 and 2008, they find significant and robust enough evidence of liquidity premia to conclude that CDS spreads cannot be used as a pure measure of credit / default risk. Their results hold for the last two quarters of 2008, when both CDS spreads and bid-ask spreads increased dramatically, arguably due to counterparty risk and deleveraging.

In addition to methodological imperfections, the Negative CDS Basis approach also suffers from impracticality, despite its easy-to-compute nature. Finding equivalent CDS contracts and corporate bonds by issuer and maturity is not straightforward. Using an index to approximate introduces further practical difficulties: CDS indices are not widely available across different economies, do not cover high yield bonds as the market for these CDS contracts is very thin and using an index can result in a mismatch with bond index constituents having substantial impact on the result. In case a (set of) bond(s) can be perfectly matched with a (pool of) CDS, ignoring the above mentioned methodological issues and assuming the CDS spread is a pure measure of credit risk, the estimated liquidity premium should not be generalized to a wider set of bonds (of similar duration, notional amount or sector). Moreover, the availability of CDS contracts is self-selecting in nature, with more illiquid bonds unlikely to have an active CDS market and any estimated liquidity premium is likely to understate the liquidity premium on bonds for which no active CDS market exists. Given the self-selecting nature the liquidity premia derived from a Negative CDS Basis approach that matches bonds and CDS directly, they can only ever be applied to the same set of bonds.

2.2.4 Structural Model Approach

Merton's model is of the structural kind, meaning it attempts to describe the explicit relationship between default probabilities, value of the firm and capital structure. Merton's model makes clever use of option pricing theory by treating a company's equity as a call option on its assets. The original model by Merton (1974) is intuitive and relatively straightforward; it assumes a very simple debt structure of one zero coupon bond maturing at time T and the underlying value of the firm follows a standard Geometric Brownian Motion;

$$dV = \mu V dt + \sigma V dz,$$

where V is the total value of the firm, μ is the asset drift, σV is the volatility of assets and dz is a standard Wiener process.

If, at time T , the value of the company's assets are lower than the amount of debt (interest and principal, D) due to be repaid, it is (in theory) rational for the company to default, leaving an equity value of zero. In case $V_T \geq D$, the company should make the repayment, leaving an equity value of DV_T . Then, from standard option pricing theory (Black and Scholes, 1973), the value of the company's equity (E) at time T is given by the pay-off structure of a call option on the value of the assets with a strike price equal to the required repayment on the debt;

$$E_T = \max(V_T - D, 0)$$

Using the standard Black and Scholes, (1973) formula to solve for the value of the equity today;

$$E_0 = V_0 \mathcal{N}(d_1) - De^{(-rT)} \mathcal{N}(d_2)$$

where

$$d_1 = \frac{\ln(\frac{V_0}{D}) + (r + (\frac{\sigma_V^2}{2})T)}{(\sigma_V \sqrt{T})}$$

and

$$d_2 = d_1 - \sigma_V \sqrt{T},$$

with constant risk-free rate r , and volatility of assets σV . The risk-neutral default probability, using a drift rate of r , of the debt is given by $\mathcal{N}(-d_2)$. In order to calculate the value of the equity, V_0 and σV , are required but unobservable quantities in the marketplace. E_0 however, is observable equity volatility (σE) can be estimated from Ito's Lemma. This gives a pair of simultaneous equations that can be solved numerically for values of V_0 and σV ;

$$\sigma_E E_0 = \frac{\partial E}{\partial V} \sigma_V V_0$$
$$\sigma_E E_0 = \mathcal{N}(d_1) \sigma_V V_0$$

Empirical applications of Merton's model show that there is some discrepancy between the default probabilities produced by Merton and those observed historically, in particular at (very) short maturities. The fair credit spreads estimated by Merton's model are reported to underestimate the observed credit spread from the market, as found by Jones et al. (1984) who are one of the first to calibrate Merton's model to real-world data. Bohn (2000), in his survey of risky debt valuation using option pricing, similarly found model spreads to be consistently lower than observed credit spreads. Merton's model has been criticised for making a set of strict assumptions that may not be reflecting the real-world accurately enough to produce realistic outcomes. As a result, Merton's model has been subject to many extensions over several decades, each extension addressing some of the simplifying assumptions made by Merton. The list of extensions below is by no means exhaustive, but includes a brief discussion of several well documented extensions/variations to model that have been applied to real world data.

- Default in Merton's model can only occur at the debt's maturity date. The model can be modified to allow for early default by introducing a threshold level so that default occurs as soon as V_t falls below this level. In some models this can be the result of shareholders' strategy to maximize equity value (Fan and Sundaresan, 2000). Models with a default barrier were pioneered by Black and Cox (1976) and are often referred to as 'first passage' models. Within first passage models, an important distinction needs to be made; those specifying an *exogenous* default boundary and those specifying an *endogenous* boundary. A typical application of a structural model with an exogenous default barrier can be found in Longstaff and Schwartz (1995). They extend the work of

Black and Cox (1976) by not having the barrier to be a (discounted) constant, but to have it evolve according to its own stochastic process. An example of a structural model that incorporates an endogenous default barrier, that is, a default barrier that is dependent on model parameters and the current state, is the model by Leland (1994). The bankruptcy point in this model is the value of the firm such that the market price of equity drops to zero. The model in Leland (1994) and Leland and Toft (1996) in particular, are scrutinised in greater detail in Section 2.8.

- The firm's assets are modelled by a Geometric Brownian Motion in Merton's original model, implying a log-normal distribution of assets. Critics have argued that a simple diffusion model might not be sufficient to capture the real world dynamics of the firm. Firms that are not currently in financial distress are, under the diffusion assumption, never in danger of defaulting on (very) short term debt obligations; i.e. firms cannot default unexpectedly. If the diffusion dynamics were accurate, (near-) zero credit spreads on short term debt would be observed, which is strongly rejected. Credit spreads on bonds with short maturities are non-zero, positive and often substantial. Both Fons (1994) and Sarig and Warga (1989) report that near-zero credit spreads on short term debt does not agree with observed credit spreads, and that the yield spread curves of certain bonds are not upward sloping as implied by the diffusion model, but are flat or even downward sloping. To allow for unexpected defaults and subsequently increase the default probability on short term debt, a jump-diffusion model for the firm's assets can be used within the Merton framework. Zhou (2001) observes that by incorporating jump risk in the default process, the model matches the size of the credit spreads more closely and the yield curve can take various shapes (upward, downward and humped), even if the firm is not in financial distress. Mason and Bhattacharya (1981) were the first include jump processes in the valuation of risky debt. Zhou (2001) claims that his model, based on a continuous diffusion process and a discontinuous jump process is more realistic compared to their model in which

the evolution on the firm follows a pure jump process with the jump amplitude following a binomial distribution, for reasons of flexibility and generality.

- The assumption of a constant risk free interest rate is not realistic and a stochastic interest rate model can be incorporated into Merton's model or any of its extensions. This also allows for the stochastic element of the evolution of the firm to be correlated with the interest rate process, if this is desired. One example of a model where interest rates are modelled by a stochastic process is a study by Shimko et al. (1993), who estimate the effect of asset and interest rate correlation on credit spread. They incorporate the short term nominal Vasicek interest rate model Vasicek (1977) into Merton's model and report that for a correlation of -0.25 between interest rate and asset dynamics, the estimated credit spread is between 5 and 7 basis point below the estimated credit spread when using a non-stochastic interest rate. Both Kim et al. (1993) and Longstaff and Schwartz (1995) show and argue that introducing a stochastic interest rate might be conceptually right, but has a relatively small effect on credit spreads while substantially increasing the complexity of the analysis. They argue that the cost of added complexity does not weigh up against introducing stochastic rates for conceptual reasons. Leland (1994), also working with a non-stochastic interest rate, notes that stochastic rates do not only unnecessarily complicate calculations, even from a real-world perspective they add very little. Referring to the criticism by Jones et al. (1984) that a contingent claims approach to valuing risky debt produces too small credit spreads, he argues that stochastic interest rates in the Merton framework lower spreads (assuming negative correlation between the asset and interest rate processes) and therefore do not 'solve the problem' of small spreads.
- Assuming a single zero-coupon bond as the total debt structure of the firm, or trying to map all existing debts into a single zero-coupon bond is clearly restrictive and not a realistic model assumption. The first structural credit risk model that allows for multiple debts with different characteristics was the Geske Compound Option model developed by Geske (1977). More recent work

has tackled the oversimplified capital structure of the firm and several models allow for more complex capital structures. Extensions have been made to include coupon paying debt; an example is Nielsen et al. (2001) who extend the basic Merton model to include an exogenous stochastic default boundary that is triggered when cash flows are unable to meet interest payments. In addition to models that include coupon bearing debt, more complex capital structure can also be modelled. Leland (1994) assumes a straightforward model where debt is perpetual and pays a continuous coupon stream, and Fan and Sundaresan (2000) build on this model by also assuming single-layered perpetual debt but including negotiations between creditors and shareholders in case of distress. In order to avoid inefficient liquidation, the model allows for shareholders to service debt strategically, with bargaining power η . For $\eta > 0$, the default barrier is lower than its counterpart in Leland (1994), with the model defaulting to the Leland (1994) when $\eta = 0$. Leland and Toft (1996) move away from the perpetual capital structure by assuming the firm continuously issues debt of the same maturity, implying the firm is also redeeming debt issued many years ago. Therefore, at any given moment in time, the firm has various debt obligations outstanding of various durations, which are all to receive coupon payments. Shorter maturities place a greater burden on the firm's cash flow because of the debt that needs to be redeemed and as a result the endogenous default barrier is much higher in the Leland and Toft (1996) model than in Leland (1994), especially for shorter specified maturities. When the maturity goes to infinity, Leland and Toft (1996) converges to Leland (1994).

- Almost all extended models include a continuous payout as part of the evolution of the firm's assets. Adjusting the drift rate to account for negative cash flow can be thought of as dividend payments to its shareholders. This simple extension was offered by Merton himself in his first publication and is analogous to the pedagogical extension of the Black-Scholes model with continuous dividend payments. In addition and similar to the inclusion of a payout ratio

in the evolution of the firm's assets, many extensions (see for example, Leland (1994) or Fan and Sundaresan (2000)), include tax advantages in their model. This ties in directly with model extensions regarding the firm's capital structure since the tax advantage is applied to coupon payments and is added to the evolution of the firm's assets as a tax shield.

- Further refinement of a more complex / realistic debt structure is needed as debt does not only vary in maturity but can have various seniority / priority structures too. Several extensions of Merton's model allow for debt to be issued with various levels of seniority (see for example, Benos and Papanastopoulos, 2007). Debt of lower seniority is valued at a higher credit spread due to a decreased recovery rate in the event of default. Recovery rates represent the percentage of the face value or market value of the debt that is received by debt holders in case of bankruptcy of the firm. These recovery rates are exogenous in structural models but are thought to vary on a firm-to-firm basis in the real world, depending on the outstanding debt across seniority classes, but also depending on the state of the economy. Recovery rates are calibrated to fit historical averages across the market. Models that take into account seniority of debt assume that absolute-priority rules are fully adhered to in case of default so that outstanding debt is paid off in strict order of seniority. Empirical evidence suggests that absolute priority rules are violated in reality; Franks and Torous (1994) investigate firms that have either entered bankruptcy through Chapter 11 of the Bankruptcy Reform Act 1978 or have informally completed a distressed exchange of traded debt. Firms entering bankruptcy under Chapter 11 only do so after attempting to resolve financial issues informally, but these firms often benefit from specific provisions in Chapter 11 because they are less solvent or liquid. As a result, the average expected recovery rate for creditors claims under a Chapter 11 bankruptcy are approximately 30% lower than recovery rates under more informal proceedings. In addition, Eberhart et al. (1990) examined bankruptcy proceedings under the same Act and measured the amount paid to shareholders in excess of

what they should have received under strict absolute priority rules. Eberhart et al. (1990) report this percentage to be 7.6% and also found evidence that equity markets expected deviations from the absolute priority rule as common share values reflect part of the value ultimately received by a violation of the strict rules.

Empirical results Structural models of default attempt to mathematically describe how default occurs, and what dynamics exist between the various factors playing a role in default or the value of the firm. Most of the structural models that relax one or more of the simplifying assumptions are theoretical in nature and describe dynamics of the firm and the dynamics of the firm's debt. No structural model tries to relax all of the simplifying assumptions nor do any of the structural models claim to be an accurate representation of the real-world dynamics of the firm. Instead, structural models allow the dynamics of how one (or several) specific factors relate to probability of default / credit spread to be studied. Examples include what effect stochastic interest rates and a correlation between firm's asset process and interest rate process has on the default probabilities or how the effect of a jump-diffusion process for the firm's assets can describe default probabilities over short maturities better. The models themselves are theoretical in nature, the credit spreads (or default probabilities) produced by these models require calibrated values for all the parameters in the model.

The standard reference for the discrepancy between credit spreads observed historically and those produced by (early) structural models after calibration is Jones et al. (1984), who reported the yield spreads derived from structural models to underestimate historical spreads. They report that on average, predicted spreads underestimate observed spreads substantially, the errors being the largest for speculative grade bonds. In addition they find that pricing errors are significantly related to maturity, equity variance and leverage. Ogden (1987), in a similar study of bond prices between 1977-1981, finds that the Merton model underestimates spreads by 104 basis points on average. Both studies emphasized the lack of stochastic interest rates in their conclusions. Given the time period under study, this is no surprise

since US Treasury rates were extremely volatile as a result of the Federal Reserve's money supply target during 1979-1982, when inflation reached double-digits. Few implementations of structural models using individual bond prices appear until Lyden and Saraniti (2000), who use a sample of non-callable bonds to calibrate and fit both Merton's original model (modified to treat coupon bonds as if they were a portfolio of zero-coupon bonds, each of which can be priced using the standard zero-coupon version of the model) and the model derived by Longstaff and Schwartz (1995), in which the firm issues a constant amount of new coupon paying debt with a fixed maturity and in which equity holders have to the option to issue new equity to service the debt or default. They find that both models underestimate yield spreads and pricing errors are correlated with coupon rate and maturity.

More recently, Eom et al. (2004) have taken five structural models, calibrated the required parameters of all models to the same sample of 182 bonds during the period 1986-1997 and compared spread predictions across models with observed historical spreads. Using the Fixed Income Database they chose a very particular set of bonds to include in their study, meeting strict requirements; only non-financial firms are included to ensure the leverage ratio is comparable across the sample, utilities are excluded since return on equity and revenues (and thereby probability of default) are dependent on regulatory influences, only fixed-rate coupon bearing bonds that are not convertible and only bonds with a simple capital structure (firms with a maximum of two publicly traded bonds and exclusion of subordinated debt) are included. In addition, the chosen firms must have publicly traded stock in order to qualify for a structural approach in the first place. Eom et al. (2004) implement five structural models (Merton, 1974; Geske, 1977; Longstaff and Schwartz, 1995; Leland and Toft, 1996 and Collin-Dufresne et al., 2001) and, contrary to previous empirical literature, fail to conclude structural models are incapable of producing sufficiently high yield spreads. They do agree that the five structural models cannot accurately price corporate debt, but the difficulties are far from limited to an underestimation of spreads. Whereas Merton (1974) and Geske (1977) are reported to consistently underestimate yield spreads, as previous work indicated, the Longstaff and Schwartz

(1995), Leland and Toft (1996) and Collin-Dufresne et al. (2001) models, on average, produce yield spreads that are too high. Both the Longstaff and Schwartz (1995) model and the Collin-Dufresne et al. (2001) model have an incredible dispersion of predicted spreads; often they are either very small or extremely large. That across the sample of bonds this averages out (slightly) higher on average is rather irrelevant as prediction error on a bond-to-bond basis is often a magnitude several times the average prediction error. The Leland and Toft (1996) model is different in the sense that it appears to consistently produce yield spreads that are too high, which Eom et al. (2004) attribute to the simplifying assumptions about coupons. The most relevant empirical application of a structural model to UK corporate bond data is in Churm and Panigirtzoglou (2005), where the Leland and Toft (1996) model is used to estimate the non-credit (liquidity) component of credit spreads. This work, published by the Bank of England, is discussed in some detail in Section 2.8

2.3 Description of the Data

The dataset used in this Chapter comes from Markit, a financial information services company providing independent data, valuations and trade processing of assets. Markit aims to enhance transparency in financial markets and improve operational efficiency for its clients. The iBoxx indices provide fixed income bond data on a daily basis, essential for many market participants in structured products, fixed income research, asset allocation and performance evaluation. iBoxx indices cover Euro, Sterling, Asian, US Dollar denominated markets, both investment grade and high yield. In addition to daily consolidated prices, a range of analytical values are provided for all the bonds in the Markit bond universe. For the research, historical iBoxx GBP Investment Grade Index data was available.

Eligibility for inclusion in the iBoxx GBP Investment Grade Index is based on several selection criteria. The following bond types are specifically excluded: bonds with American call options, floating-rate notes and other fixed to floater bonds, optionally and mandatory convertible bonds, subordinated bank or insurance debt with mandatory contingent conversion features, CDOs or bonds collateralized by

CDOs. In addition, retail bonds and private placements are reviewed by the iBoxx Technical Committee on an individual basis and excluded if deemed unsuitable (Markit, 2012a).

All bonds in the Markit iBoxx GBP universe must have a Markit iBoxx Rating of investment grade (Markit, 2012b). The average rating of Fitch Ratings, Moody's Investors Service and Standard & Poor's Rating Services determines the iBoxx rating. Investment grade is defined as BBB- or higher from Fitch and Standard & Poor's and Baa3 or higher from Moody's. Ratings from the rating agencies are converted to numerical scores and averaged, then consolidated to the nearest rating grade; the iBoxx Rating system does not use tranches. Eligibility for inclusion is also conditional on the amount outstanding, where the issue needs to be of a minimum size. Gilts need to have an outstanding amount of at least GBP 2bn, whereas the minimum amount for non-Gilts is set to 250m.

The sample (Oct 2003 - Jul 2014) of 2767 trading days includes 2392 unique bonds from 744 different issuers (by Ticker), with data for approximately 900 bonds on any given day. The analysis uses a range of analytical values (Markit, 2014) included in the index. These bond characteristics can be contractual (e.g. coupon rate, issuer, maturity, seniority, date of issue, industry) or time dependent (e.g. bid- and ask prices, credit rating, credit spread).

All bonds are classified based on the principal activities of the issuer and the main sources of the cash flows used to pay coupons and redemptions. In addition, a bond's specific collateral type or legal provisions are evaluated. Hence, it is possible that bonds issued from different subsidiaries of the same issuer carry different classifications. The issuer classification is reviewed regularly based on updated information received by Markit, and status changes are included in the indices at the next rebalancing if necessary. This implies that even though several analytical values or variables in the dataset are considered 'static', there is the possibility of them changing. The main sector classifications within the Markit iBoxx GBP Index are the following:

- *Gilts*. Bonds issued by the UK central government denominated in Sterling.

- *Sovereigns*. Bonds issued by a central government other than the UK and denominated in Sterling.
- *Sub-Sovereigns*. Bonds issued by entities with explicit or implicit government backing due to legal provision or the public service nature of their business.
 - *Agencies*. Bonds issued by entities whose major business is to fulfil a government-sponsored role to provide public, non-competitive. Often, such business scope is defined by a specific law, or the issuer is explicitly backed by the government.
 - *Supra-nationals*. Bonds issued by supranational entities, i.e. entities that are owned by more than one central government (e.g. World Bank, EIB).
 - *Public Banks*. Bonds issued by publicly owned and backed banks that provide regular commercial banking services.
 - *Regions*. Bonds issued by local governments (e.g. Isle of Man)
- *Corporates*. Bonds issued by public or private corporations. Bonds secured by a floating charge over some or all assets of the issuer are considered corporate bonds. Corporate bonds are further classified into Financials and Non-Financials bonds and then into their multiple-level economic sectors, according to the issuer's business scope. The category insurance-wrapped is added under Financials for corporate bonds whose coupon and principal payments are guaranteed by a special mono-line insurer.

In addition to a classification based on the issuer's business scope and activities, corporate debt is further classified into senior and subordinated debt. Subordinated debt is mostly issued by financials, but other corporate issuers might be forced to do so if indentures on earlier issues mandate their status as senior bonds. Subordinated debt can be especially risk-sensitive since the bond holders only have claims on an issuer's assets after other bond holders, but without the upside potential that shareholders enjoy. Capital in the form of debt instruments is always sub-ordinated because senior debt does not count towards bank capital. From a regulatory point

of view the bank's capital serves as protection of depositors, a safety net that can absorb unexpected losses to guarantee depositors. A financial institution's debt can be categorized as one of the following:

- *Tier 1.* Shareholder's equity and retained earning are commonly referred to a banks 'core' (Tier 1) capital for regulatory purposes. The Tier 1 capital, as issued debt, consists of other securities qualifying as Tier 1.
- *Upper Tier 2* From a regulatory perspective of a bank's capital, Tier 2 debt comprises undisclosed reserves, revaluation reserves, general provisions, hybrid instruments (preferred) and subordinated term debt.
- *Lower Tier 2* From a regulatory perspective of a bank's capital, only 25% of a bank's total capital can be Lower Tier 2 debt; it is easy and cheap to issue. In order to ensure that a bank's capital from subordinated debt issues does not fall substantially after and issue matures, the regulator demands that Tier 2 capital debt amortises on a straight line basis from maturity minus five years.

The market information on the tier of subordination for insurance capital is often less standardized than the equivalent issues by banks. In these cases, the classification is based on the maturity, coupon payment and deferral provisions of the bond from the offering circulars of the bonds (Markit, 2012a).

Bonds with option-like characteristics (embedded options) are included in the dataset. These option-like features are part of the bond rather than separately traded securities and the option-like characteristics are not necessarily mutually exclusive; one bond may have multiple option features embedded. Embedded bonds can include but are not limited to:

- *Callable Bonds.* Bonds that give the issuer the option of buying back the bond at a predetermined price at some point in the future. The lockout period refers to the initial time period in which the bond cannot be redeemed by the issuer.
- *Puttable Bonds.* Bonds that give the bond holder the option to demand early redemption at a predetermined price at some point in the future.

- *Convertible Bonds.* Bonds that give the bond holder the option to demand conversion of bonds into stocks at a predetermined price at some point in the future.

Depending on the type of option feature(s) embedded in a bond, the Credit Spread can either be higher or lower than the Credit Spread for an equivalent bond without option features. No information about a bond's optionality is included in the dataset explicitly.

The database reports a number of different measures of the credit spread including the annual benchmark spread (ABS), Option-Adjusted Spread and Z-spread (Markit, 2014). Of these, only the annual benchmark spread is reported for the full duration of the dataset and this is used as the measure of the credit spread in all statistical analyses. The analyses have been repeated over shorter periods with the alternative measures of credit spread and the results are found to be robust (results appear in Section 2.7.2).

Markit iBoxx index calculations are based on multi-sourced pricing which, depending on the structure of each market, takes into account a variety of data inputs such as transaction data, quotes from market makers and other observable data points. For the GBP Corporate Index, the source of data is quotes from market makers. Currently ten market makers submit prices, including Barclays Capital, Goldman Sachs, HSBC, Deutsche Bank and JP Morgan. All submitted prices and quotes have to pass through a three-step consolidation process before being included in the end-of-day value (Markit, 2008).

In addition to descriptive classifications of a bond, a range of analytical values is available on a daily basis. This section aims to give an overview of the most important of these analytical values, focusing on Credit Spread measures in particular. The analytical values that are available on a daily basis include key measures such as Years to Maturity, (Modified) Duration, (Annual) Yield, (Annual) Convexity and Age of the bond. In addition to these key measures, many more are reported on a daily basis including Daily Returns, Month-Date>Returns, Number of Contributors, Excess Return over Sovereigns, Duration weighted exposure; all yield calculations

use the bid price of the bond.

During the course of the sample period of the study (October 2003 – July 2014) there have been more than a handful of changes to the availability of certain analytical values. Most importantly, the available pieces of information of every bond on any given day has increased from twenty in October 2003 to more than forty in July 2014. For this reason, to maintain backwards compatibility in order to include the entire sampling period, many of the newer bits of information need to be excluded. This is not a problem for the majority of information as they add very little value (for example, seven buckets of Expected Remaining Life), but when it comes to Spread calculations it is important to note what changes have occurred during the sample period. For the last day included in the sample, four different credit spread measures are reported on a daily basis.

The only spread measure that is available for the entire period under study is the Annual Benchmark Spread, mathematically;

$$\text{BMS}_{i,t}^a = Y_{i,t}^a - Y_{BM(i),t}^a$$

where $Y_{i,t}^a$ is the annualized yield of bond i at time t and $Y_{BM(i),t}^a$ is the annualized yield of benchmark bond i at time t . Note that for benchmark bonds $\text{BMS}_{i,t}^a = 0$. On any given day 20-25 Gilts are classified benchmark bonds. For example, on 9th of October 2012, 23 Gilts are classified as benchmark bonds and used to calculate the Annual Benchmark Spread for all bonds in the dataset on that particular day. Figure 2.1 plots those 23 bonds (years to maturity versus annual yield), and illustrates the credit spread.

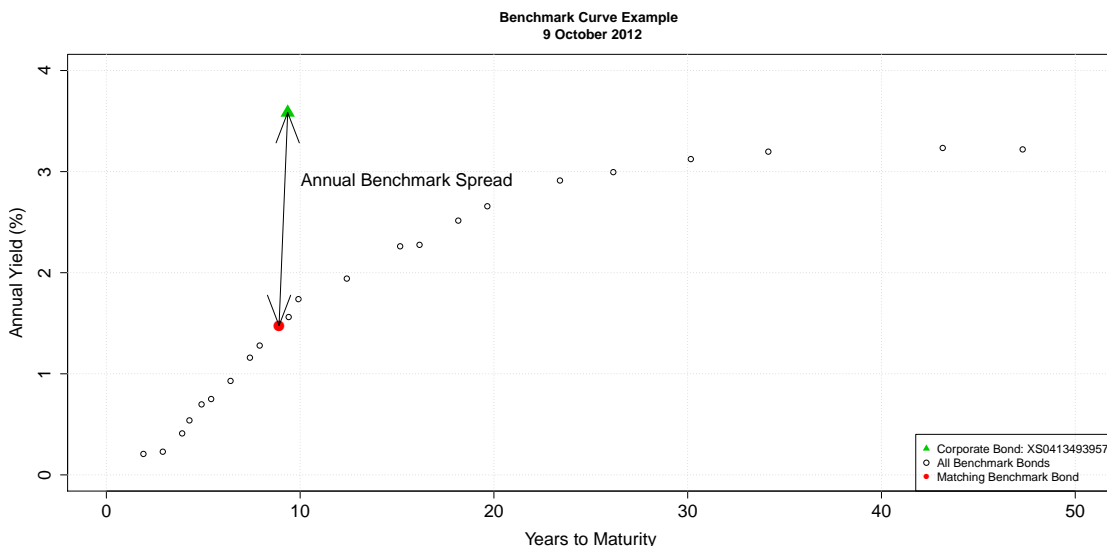


Figure 2.1: The Annual Benchmark Spread is based on the difference in annual yield between a bond and the corresponding benchmark bond. A benchmark bond is chosen as to minimise the difference in maturity (remaining life) between a bond and its reference.

It is important to mention that the methodology of assigning a specific benchmark bond to each of the bonds in the dataset could, potentially lead to jumps in credit spread that are purely ‘mechanical’. The maturity of the bond and its assigned benchmark bond are expected to decline in tandem; if, however, a bond were to switch benchmark bond there could be a jump/fall in credit spread due to a change of the reference bond. Since the anticipated effect of this phenomenon is negligible, the Annual Benchmark Spread is used directly.

2.4 Descriptive Analysis of Credit Spreads

The Markit iBoxx dataset provides Annual Benchmark Spreads, for all bonds on every trading day which makes computations readily available. This allows Credit Spreads to be monitored through time for all individual bonds, but also allows the description of Credit Spreads at more aggregate levels; one can think for example, of Spreads aggregated by Rating, by Financials and Non-Financials, by Senior and Subordinate debt, or by a combination of several of the above. When examining data by means of plots, the data will be split in eight categories or groups; split by Rating (AAA, AA, A and BBB) and Financials / Non-Financials. The split into eight categories is granular enough to get a good idea of how different segments of

the market behaved at any one point, yet sufficiently aggregate not to get lost in details.

To describe the distribution of Credit Spreads of the eight categories defined above on any given day during the sample period (Oct. 2003–Jul 2014), the Credit Spread time series is plotted using the mean Credit Spread as well as the interval between which 90% of all observed Credit Spreads lie;




Figure 2.2: Daily average Credit Spreads (and 5th and 95th quantile) over time for Financial/Non-Financials and the four investment grade rating categories.

Several, perhaps obvious, observations can be made from the time series plots in Figure 2.2:

- Credit Spreads move together, regardless of Rating and Financial status of the issuer. Credit Spreads across all eight groups start to increase mid-2007, peak early 2009 and decline rapidly in the following year.
- Credit Spreads for Financials and Non-Financials seem to be at very similar levels (within a given Rating category) in the years prior to the ‘credit crunch’. During the crisis years, spreads of Financials rise substantially, far more than spreads on Non-Financials; this is no surprise given the nature of the crisis.
- Credit Spreads increased in 2011 in response to the sovereign debt crisis in Europe. The Credit Spreads peaked early 2012 and have since been declining. The downward trend continues until the last day in the dataset in July 2014.

2.4.1 Timeline Analysis

Rather than treating ‘the credit crunch’ and the subsequent ‘European sovereign crisis’ as single events or single periods of time, major financial events can be described in isolation. This is used to subjectively quantify the impact certain events had on Credit Spreads of each of the eight groups (split by rating and financial status) and perhaps whether the market responded prior to the event occurring, indicative of market expectations. Sixteen major financial events have been selected and can be seen in Figure 2.3.

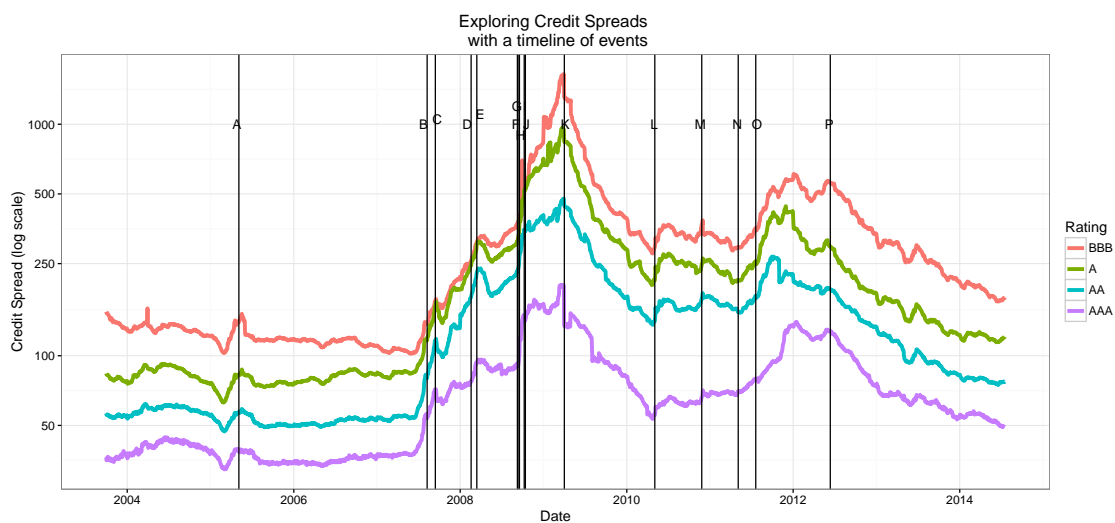


Figure 2.3: Average Credit Spread by Rating (on a logged y-axis) over time, annotated with a non-exhaustive selection of events during the ‘credit crunch’. Tick marks indicate the start of the year.

- **A, 5th May 2005.** Tony Blair is re-elected as Prime Minister of the United Kingdom by popular vote, with 35.2%, the lowest majority government in British history.
- **B, 9th August 2007.** BNP Paribas indicate they cannot value the complex assets (CDOs) for three of their funds, for which trading freezes. It is the first major bank to acknowledge the risk of exposure to sub-prime mortgage markets. Northern Rock’s chief executive later reflects saying that it was ‘the day the world changed’.
- **C, 14nd September 2007.** Northern Rock has borrowed large sums of money to fund mortgages for customers, and needs to pay off its debt by re-

selling those mortgages in the capital markets. Due to fallen demand, Northern Rock faces a liquidity crisis and it needs a loan from the British government. This sparks fears that the bank will shortly go bankrupt; Britain's first bank run since 150 years was a result.

- **D, 17nd February 2008.** The government nationalises the troubled mortgage lender Northern Rock as a result of prolonged liquidity problems.
- **E, 14th March 2008.** The investment bank Bear Stearns is bought out by JP Morgan. This makes Bear Stearns the biggest player fallen in the crisis.
- **F, 7th September 2008.** The US government bails out Fannie Mae and Freddie Mac; two huge firms that had guaranteed thousands of sub-prime mortgages.
- **G, 15th September 2008.** Lehman Brothers, deeply involved in the sub-prime mortgage markets, files for bankruptcy causing worldwide financial panic.
- **H, 17th September 2008.** The UK's largest mortgage lender, HBOS, is rescued by Lloyds TSB after a huge drop in its share price.
- **I, 8th October 2008.** Iceland's three biggest commercial banks Glitnir, Kaupthing, and Landsbanki collapse. To protect the deposits of their many British customers, Gordon Brown uses anti-terror legislation to freeze the assets of the banks' UK subsidiaries.
- **J, 13th October 2008.** The British government bails out several banks, including the Royal Bank of Scotland, Lloyds TSB, and HBOS.
- **K, 2nd April 2009.** The G20 agrees on a global stimulus package worth \$5 trillion.
- **L, 2nd May 2010.** Signalling the start of the Eurozone crisis, Greece is bailed out for the first time, after Eurozone finance ministers agree loans worth 110

billion Euros. This intensifies the austerity programme in the country, and sends hundreds of thousands of protesters to the streets.

- **M, 28th November 2010.** European ministers agree a bailout for Ireland worth 85 billion Euros.
- **N, 5th May 2011.** The European Central Bank bails out Portugal.
- **O, 21st July 2011.** Greece is bailed out for a second time, after it failed to get all its affairs in order.
- **P, 6th June 2012.** The level of Spanish borrowing reaches a record high, indicating the Euro-crisis is still on-going and recovery is a slow process.

Prior to the May 2005 parliamentary elections (event A) Credit Spreads had been increasing for some time. Elections were announced on April 4th, a month before the elections on May 5th, but the general election had been covered in the media for weeks. Uncertainty over the outcome and the lead up to the election may have spurred higher spreads. The subsequent two events, B and C, are often regarded as the events that signify the start of the financial crisis and coincide with the start of the increase in Credit Spreads throughout the crisis. Looking at the time series closely is clear that the increase in spreads precedes the actual events, indicating the market seemed to perceive an increased risk several days/weeks before the day labelled as the event. Looking very closely at B and C it can be seen that immediately after Northern Rock's loan from the UK Government (event C), the Credit Spread appears to drop, or at least move sideways for a brief period of time. The unrest in September 2008 later appeared to be a surprise to the market as the large jump in spreads exactly coincides with the events. One could argue that for financial markets this might have been the true start of the crisis, given the tremendous jump in spreads. Spreads continue to increase until the UK government bailout of the Royal Bank of Scotland, Lloyds TBS and HBOS on 13 October 2008. Whereas the bailout coincides with the peak observed for Financials, it appears as if the non-Financials peaked a few months earlier, and were on the way to recovery before Financials. The bailout seems to start a steady downward trend

which continues well into 2011, only to be interrupted for a very brief period of time due to the bailing out of Greece (event L) and Ireland (event M). The bailing out of Portugal followed by the second bailing out of Greece two months later (events N and O) caused another steady increase in Spreads until 2012 when again Spreads were slowly declining, a trend that is slow and continues until the last day in the dataset.

2.4.2 Descriptive Analysis of Market Data

The dynamics of movement in Credit Spreads over time have been described in detail in the previous section. This section aims to provide further understanding of, on the one hand, the available pieces of information contained in the iBoxx dataset, and on the other hand aims to provide some insight into the dynamics of the corporate bond market it describes. Firstly, the amount of data is investigated in Figure 2.4 using the number of bonds in the dataset (top left), the total notional amount of debt outstanding (top right), the average duration of bonds (bottom left) and the number of ‘new’ bonds (bottom right).

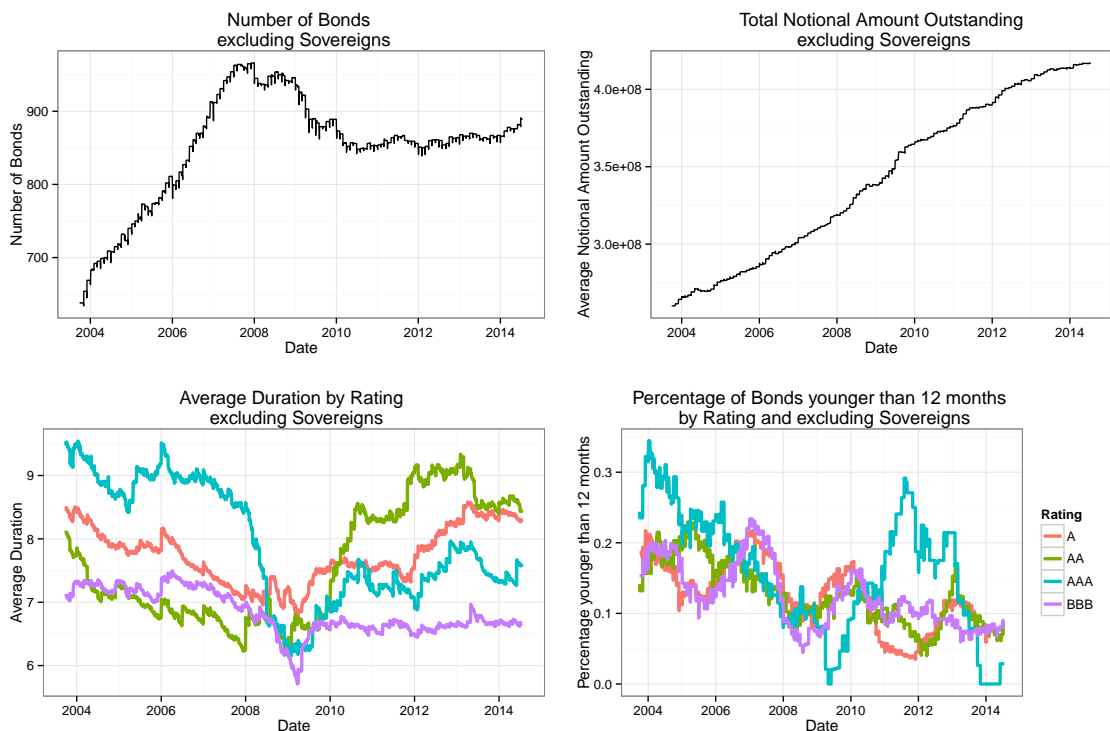


Figure 2.4: Descriptive analysis of market data: number of issuers (top left), total notional amount outstanding (top right), average duration by rating (bottom left) and the proportion of the bonds universe issued fewer than twelve months ago.)

Based on the two graphs at the top of Figure 2.4, one can see that the number of unique issues (investment grade only) peaked around 2008, when, possibly in response to ongoing financial distress, firms were likely to postpone new issues. The amount of outstanding debt however, has increased over the entire period under study. Figure 2.4 (bottom left) shows how the average duration of the observed bonds varies over time. In general, the figure shows how there appears to be a response to the financial distress leading to a decline in observed durations; this is certainly the case for AAA, A, and BBB-rated bonds, whereas the effect of AA-rated bonds is more difficult to observe. Figure 2.4 (bottom right) gives an indication of the issuance of new bonds, as it shows the percentage of bonds that were issued less than twelve months ago. One can easily observe that across rating categories, the numbers are very similar, with the post-2009 numbers for AAA-rated bonds being very volatile due to the number of bonds in this rating category in general.

From a very high level, Figure 2.5 shows how the market for investment grade corporate bonds has changed during the period under study. For instance, the average Bid-Ask Spread of bonds (Figure 2.5, left) evolves over time and bivariate relationships (based on correlation estimates) have changed (Figure 2.5, centre and left).

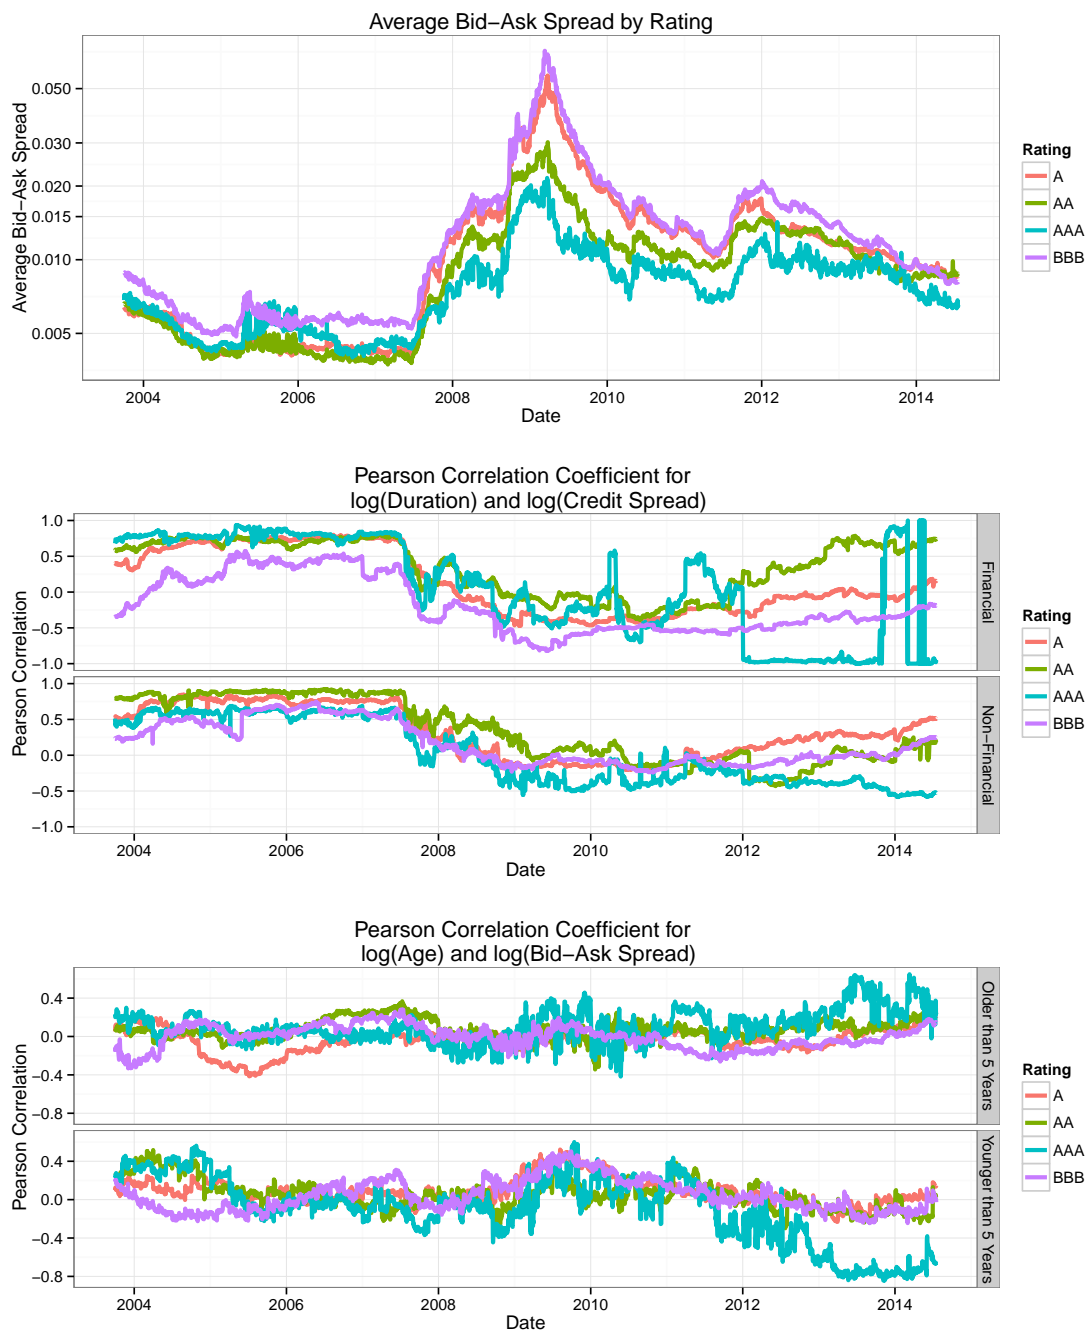


Figure 2.5: Exploratory Analysis of a selection of available analytical values and their change over time (top), coupled with the way in which some bivariate relationships vary over time (centre, bottom).

Figure 2.5 (top) provides a very high-level overview of a key component to the study of liquidity; Bid-Ask Spreads. Important to note is the log-scale of the y-axis which illustrates the extreme increase in compensation demanded by market makers to take the opposite side of trades. This appears to follow anecdotal evidence that liquidity becomes a major issue when financial uncertainty (the credit crunch) hits. Without trying to over-analyse a simplistic visualization, it does, interestingly ap-

pear as if the A-rated bonds were hit the hardest when looking at pre- and post-crisis numbers for the Bid-Ask Spread. In addition to aggregating over many underlying variables, Figure 2.5 also fails to show the width of the range of observed Bid-Ask Spreads within a rating category. The centre and bottom plot in Figure 2.5 both show Pearson correlation coefficients as a straightforward means to quantifying the bivariate relationship of several key variables. Important to note that no causal relationship is implied to exist, the plots merely show a co-occurrence of values and how this may have changed over time; for simplicity, standard errors of the correlation estimates are omitted. As such, careful to draw conclusions, it is worth noting that;

- The correlation estimates for AAA-rated bonds, particularly for those issued by Financials, are volatile due to the small number of bonds that fall within this category.
- The correlation between the Duration and Credit Spread is positive on the whole, which would be indicative of an upward sloping Duration-Spread curve, subject to controlling for non-Duration differences that cause variation in spreads.
- The correlation is highly positive prior to the onset of financial distress in 2007, when the coefficients drop substantially, across ratings and across Financial and non-Financial firms. From the end of 2010 onwards, the coefficient has been increasing, in general, until the last day in the dataset.
- Comparing across rating categories, a few observations stand out; BBB-rated Financial firms appear to have a consistently lower correlation over the entire time period, where all the other rating categories move together, with the exception of the period post-2011. After 2011, all the coefficients for Financial firms take a different path; AA recovers to pre-crisis levels very quickly, A recovers a lot more slowly even though pre-crisis coefficients were near-identical to the AA-coefficients and the BBB coefficient stays flat until a slight recovery is visible in 2013. Considering the non-Financials, coefficients are all rising together with the exception of AAA-rated bonds, for which the coefficient is

below pre-2008 levels. The caveat to this descriptive analysis is the effect all other variables have on the credit spread; these attributes may have changed over time and may have changed across rating category.

Figure 2.5 (bottom) shows the bivariate relationship of the age of a bond (measured in years since issuance) and the Bid-Ask Spread, where the universe of bonds is divided into two groups, bonds younger than five years and those older than five years. Previous literature by Houweling et al. (2005) suggests that age itself can be a good proxy of liquidity and suggests that corporate bonds may be (relatively) actively traded when they are first issued, but after some time a very large portion of the market will find its way into institutional portfolios. Based on empirical work in Houweling et al. (2005), a simple cut-off at six or twelve months may capture this effect best; hence the inclusion of this indicator variable in later analyses.

2.5 Modelling Process

The liquidity premium is defined as the difference in spread between a bond's observed spread in the market and the spread of a hypothetical bond, identical in all aspects, but perfectly liquid. Figure 2.6 illustrates this concept further (highly stylised);

- **A** represents the yield curve for risk free, perfectly liquid bonds (bid price = ask price), against which *theoretical* credit spreads are measured.
- **B** (not observable) adds in expected default losses for perfectly liquid corporate bonds of a given rating.
- **C** (not observable) adds in a risk premium for default losses (sometimes referred to as the allowance for unexpected default losses) for perfectly liquid bonds.
- **D1** and **D2** represent the ask and bid yields respectively on bonds with medium levels of illiquidity.

- **E1** and **E2** represent the ask and bid yields respectively on bonds with high levels of illiquidity.

Market credit spreads are based on bid prices (Markit, 2008). The liquidity premium is defined as the difference between an individual bond's credit spread (e.g. E2) and the credit spread for an equivalent but perfectly liquid bond (curve C). The challenge is that curve C cannot be observed and needs to be estimated.




Figure 2.6: Stylised representation of bond yields, illustrating the challenge to estimate yield curve C in order to extract liquidity premia. Yield curves: A- risk free (e.g. gilts); B- as A plus expected default losses; C- as B plus credit risk premium; D_{1,2}- as C plus liquidity premium and bid/ask spread; E_{1,2}- as D but higher bid/ask spread.

2.5.1 Modelling Methodology

To extract liquidity premia from corporate bond prices, a three stage modelling process is constructed. In the first stage the Bid-Ask Spread is modelled and a new liquidity proxy, the Relative Bid-Ask Spread (RBAS), is derived. The RBAS is a measure of a bond's illiquidity relative to bonds with identical characteristics (on the same day) and is used in the second stage of the modelling process. In

the second stage, Credit Spread is modelled as a function of bond characteristics, including the bond's RBAS. The third and final stage extracts liquidity premia by computing the difference between a bond's observed spread with the hypothetical spread on a perfectly liquid equivalent bond, estimated by extrapolation.

2.5.2 Modelling the Bid-Ask Spread

In the first stage the Bid-Ask Spread is modelled using bond characteristics. Separate cross-sectional regression models are fitted to each trading day (t), for each rating (r). A total of 2767 days \times four ratings (AAA, AA, A and BBB) means a total of 11068 regression models are fitted.

$$\begin{aligned}
 BAS(i, r, t) &= (\text{Ask Price} - \text{Bid Price}) / \text{Bid Price} \in (0, \infty) \\
 I_X(r, t) &= \text{indicator X: 0 or 1} \\
 \log(BAS(i, r, t)) &= c(r, t) \\
 &+ \beta_{1,FIN}(r, t) \times \log \text{Duration}(i, t) \times I_{FIN}(i) \\
 &+ \beta_{1,NF}(r, t) \times \log \text{Duration}(i, t) \times I_{NF}(i) \tag{2.1} \\
 &+ \beta_2(r, t) \times \log \text{Notional}(i, t) \\
 &+ \beta_3(r, t) \times \text{Coupon}(i, t) \\
 &+ \sum_k \beta_k(r, t) \times I_k(i, t) \\
 &+ \epsilon_{BAS}(i, t) \quad (\text{residual}),
 \end{aligned}$$

where $c(r, t)$ is a constant for rating r at time t , indicator variables ($I_k(i, t)$) are Financial (FIN) or Non-Financial (NF) Issuer, Sovereign or Non-Sovereign Issuer (for AAA and AA-rated bonds), Senior or Subordinate (for A and BBB-rated bonds), Collateralized or Not-Collateralized, Bond Age (Age < 1 / Age > 1) and Debt Tier (Lower Tier 2, for A and BBB-rated bonds).

The inclusion of covariates is based on both economic intuition and previous literature; Houweling et al. (2005) for example, examine the use of Issue Size, Duration and Bond Age as liquidity proxies in their regression models. Parameters

are estimated using least squares regression (OLS).

The relative liquidity measure (Relative Bid-Ask Spread) is defined as;

$$RBAS(i, t) = \exp(\epsilon_{BAS,i,t}).$$

By design, $\log(RBAS)$ is uncorrelated with any covariates included in Equation (2.1), which makes for an attractive property; RBAS can be interpreted as the bond's liquidity, independent of any bond characteristics (included as covariates). By design, the distribution of $\log(RBAS)$, for a given rating and day, is centred around 0, irrespective of rating, day or economic climate. The variance of the distribution is directly related to the quality of fit of the regression analysis in equation (2.1) and determines the variation of observed values for RBAS and ultimately contributes to the variation of the estimated liquidity premium.

2.5.3 Modelling the Credit Spread

Credit Spreads are modelled using the same approach; bond characteristics are used to explain variation in Credit Spreads, cross-sectionally, for each trading day and rating (approximately 11,000 regressions).

$$\begin{aligned} \log(CS(i, r, t)) = & c(r, t) \\ & + \gamma_{1,FIN}(r, t) \times \log \text{Duration}(i, t) \times I_{FIN}(i) \\ & + \gamma_{1,NF}(r, t) \times \log \text{Duration}(i, t) \times I_{NF}(i) \\ & + \gamma_2(r, t) \times \log \text{Notional}(i, t) \\ & + \gamma_3(r, t) \times \text{Coupon}(i, t) \\ & + \gamma_4(r, t) \times \text{RBAS}(i, t) \\ & + \sum_k \gamma_k(r, t) \times I_k(i, t) \\ & + \epsilon_{CS}(i, t) \quad (\text{residual}), \end{aligned} \tag{2.2}$$

where the indicator variables are identical to Equation (2.1).

Corporate debt is classified into senior and subordinated debt, where subordinated debt is mostly issued by financials, but other corporate issuers might be forced

to do so if indentures on earlier issues mandate their status as senior bonds. Subordinated debt can be especially risk-sensitive since the bond holders only have claims on an issuer's assets after other bond holders (without the upside potential that shareholders enjoy).

Estimated regression coefficients from equation (2.2) give an insight into which bond characteristics influence credit spreads cross-sectionally, and how this changes over time. It also allows the testing whether illiquidity is positively priced ($\gamma_4(r, t) > 0$), and whether the price of relative illiquidity varies over time (and by rating).

Whereas the linear model specification in equation (2.2) is unable to capture potential higher order effects and is restricted to the functional form imposed on the regression, it does allow for intuitive interpretation of the model coefficients (Section 2.6.2). Even though the functional form and choice of co-variates introduce a considerable amount of subjectivity, the calibration of those parameters is pre-determined and does not require any judgement on part of the modeller.

2.5.4 Credit Spread of Perfectly Liquid Bonds

As conceptually illustrated in Figure 2.6, the liquidity premium is interpreted as additional spread of an illiquid bond over its perfectly liquid equivalent, where the perfectly liquid equivalent is not observable in the market. Using regression Equation (2.2), a model is formulated to estimate the spread of a perfectly liquid equivalent bond by extrapolating RBAS to zero. Since RBAS¹ is designed to be uncorrelated with any other covariate in Equation (2.2), one can extrapolate to zero, without

¹Strictly speaking, $\log(\text{RBAS})$ is uncorrelated with the bond characteristics included in the regression modelling bid-ask spread.

having to make adjustments to other covariates;

$$\begin{aligned}
 \log(\tilde{C}S_{liq}(i, r, t)) &= \hat{c}(r, t) \\
 &+ \hat{\gamma}_{1,FIN}(r, t) \times \log \text{Duration}(i, t) \times I_{FIN}(i) \\
 &+ \hat{\gamma}_{1,NF}(r, t) \times \log \text{Duration}(i, t) \times I_{NF}(i) \\
 &+ \hat{\gamma}_2(r, t) \times \log \text{Notional}(i, t) \\
 &+ \hat{\gamma}_3(r, t) \times \text{Coupon}(i, t) \\
 &+ \hat{\gamma}_4(\mathbf{r}, \mathbf{t}) \times \mathbf{0} \quad (\text{perfectly liquid}) \\
 &+ \sum_k \hat{\gamma}_k(r, t) \times I_k(i, t) \\
 &+ \hat{\epsilon}_{CS}(\mathbf{i}, \mathbf{t}) \times \mathbf{0} \quad (\text{no residual}).
 \end{aligned} \tag{2.3}$$

Then, the Liquidity Premium (i, r, t) is easily derived from both the fitted Credit Spreads $(\tilde{C}S(i, r, t))$ and the estimated perfectly liquid equivalent Credit Spreads $(\tilde{C}S_{liq}(i, r, t))$;

$$\begin{aligned}
 \tilde{L}P_{bps}(i, r, t) &= \tilde{C}S(i, r, t) - \tilde{C}S_{liq}(i, r, t) \\
 \tilde{L}P_{\%}(i, r, t) &= \frac{\tilde{C}S(i, r, t) - \tilde{C}S_{liq}(i, r, t)}{\tilde{C}S(i, r, t)}.
 \end{aligned}$$

2.6 Numerical Results

By way of example, much of the discussion of results focuses on A-rated bonds. For other rating classes similar results are obtained and interpretation can follow the same lines. Where appropriate, charts are provided charts to compare results across ratings and primary model outcomes, such as estimates liquidity premia, are reproduced across all rating categories in Appendix A.

2.6.1 Modelling Bid-Ask Spread

Since the model in Equation (2.1) has been fitted over a relatively long period of time, the robustness of the model parameters, β_k , can be investigated over time, where robustness is defined as stability over short periods of time. Given the significant

shock financial markets endured during the credit crunch, it is reasonable to expect relationships to (temporarily) change, or at least respond, as a result. The evolution of several model parameters is shown in Figures 2.7 and 2.8. To aid the interpretation of model parameters in both the Bid-Ask Spread and Credit Spread models, it is important to note the log transformation on some variables (e.g. duration and notional amount). Whereas for indicator variables the range of possible values is well-defined (either 0 or 1), the range of possible values for log-transformed variables is less straightforward; hence, a measure of dispersion for each regression co-variate is provided in Table 1.

Rating	Duration	Not. Amount	Coupon	Age	Financial	Senior	Collateralized	Tier LT2
AAA	0.82	0.78	1.28	0.15	0.90	0.10	0.26	NA
AA	0.56	0.73	1.41	0.13	0.49	0.45	0.22	0.11
A	0.54	0.63	1.29	0.12	0.47	0.48	0.10	0.14
BBB	0.48	0.57	1.23	0.12	0.67	0.48	0.12	0.08

Table 2.2: Average of daily standard deviations by variable and rating class. Note that Duration and Notional Amount are logged variables in the model.




Figure 2.7: Beta parameters (β) for log duration, grouped by Financials and Non-Financials.

The duration beta parameter for both Financial issuers and Non-Financial is-

suers, for A-rated bonds can be seen in Figure 2.7. The duration coefficient, β_1 , is close to one (approximately 0.9) for both Financial and Non-Financial issuers, prior to the crisis. In 2008, just after the nationalisation of Northern Rock, the coefficient dropped substantially. The duration beta coefficient for Non-Financials recovered to pre-crisis levels far more quickly than its Financials counterpart (mid-2010 versus beginning 2013).

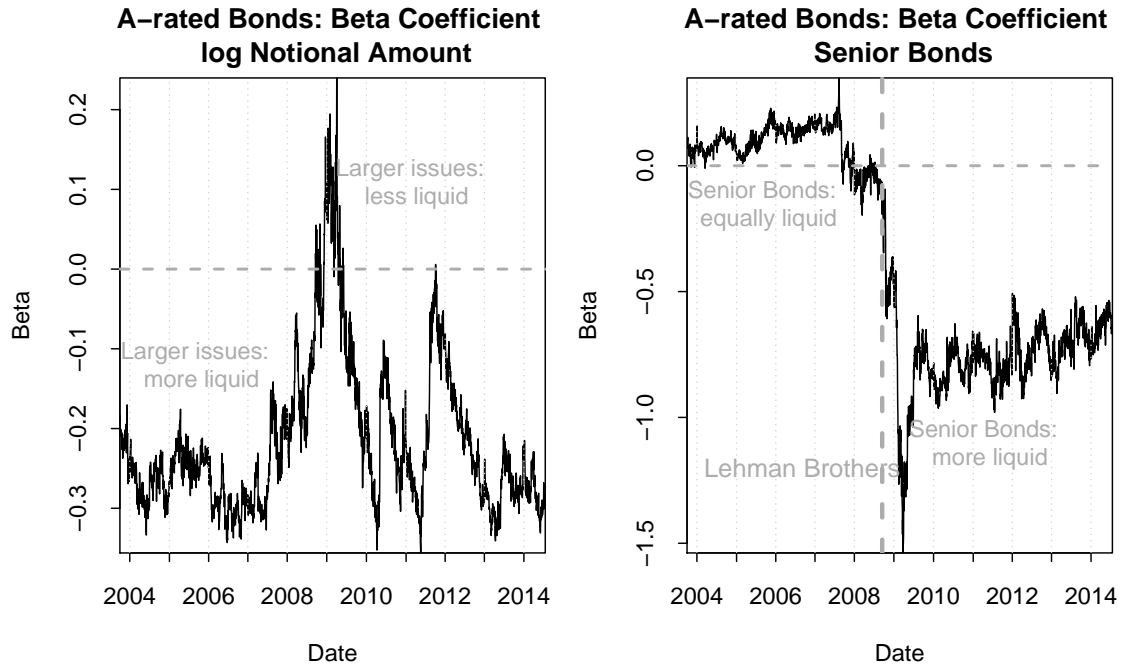


Figure 2.8: Beta parameters (β) for log Notional (left) and Seniority indicator (right).

Figure 2.8 shows the evolution of the beta-coefficient for log Notional (Amount) (left) and for the Senior / Subordinate indicator (right). The negative coefficient of the Notional Amount beta parameter (Figure 2.8, left) indicates that larger issues generally have lower bid-ask spreads. This relationship broke down at the height of the crisis in 2009, suggesting that large issues were more difficult to trade at the desired volumes. The apparent unexpected result could be a data artefact; since large issues were the only bonds trading at the time, the quotes for small bonds may not have been updated. Figure 2.8 (right) shows that Senior bonds did not trade at different levels of liquidity prior to the credit crunch. The onset of the credit crunch caused Senior bonds to trade at much lower bid-ask spreads. The increased (decreased) liquidity of Senior (Subordinate) bonds seems to support the

often-quoted ‘flight-to-quality’ of safer Senior (Financial) bonds.




Figure 2.9: Beta parameters for the Age Category indicator (left) and Capital Tier (LT2) indicator (right).

Lastly, Figure 2.9 displays the beta parameters for the indicator variables related to Bond Age (>1 year) (left) and Capital Tier (LT2) (right). The coefficient for the Age indicator (Figure 2.9, left) is rather volatile and relatively small in magnitude; with a value of approximately 0.1 on average, implying that recent issues (age < 1 year) typically have a bid-ask spread that is 10% narrower than older issues (age > 1 year). The sign of the beta parameter is according to expectations on most days; older issues (>1 year) appear to be less liquid. However, from late 2007 to the end of 2008, the coefficient was negative indicating that newly issued bonds were more difficult to trade at that time, perhaps because short-term traders suddenly found it difficult to offload recently purchased new issues. A financial institution’s debt is capital that serves as protection of depositors from a regulatory viewpoint and the regulator categorizes this capital in tiers. From a regulatory perspective of a bank’s capital, only 25% of a bank’s total capital can be Lower Tier 2 debt and is generally the easier and cheapest to issue. Not unsurprisingly, the ‘flight-to-quality/liquidity’ can also be observed in Figure 2.9 (right), where LT2 capital becomes more illiquid after the onset of the credit crunch, with extreme levels of illiquidity during early

2009.

Throughout the results section, the focus will be on displaying numerical results for the models fitted to A-rated bonds, as this corresponds most closely to the typical credit quality of an insurance portfolio. However, the modelling approach also allows for beta coefficients to be compared across rating, providing insight into market behaviour at different segments of the credit quality spectrum. For example, the duration coefficients (β_1) for Financials and Non-Financials can be compared across rating.




Figure 2.10: The weekly duration coefficient, β , for Non-Financial Issuers (left) and Financials (right) for all four rating categories.

The weekly duration coefficient for Non-Financial Issuers (Figure 2.10, left) is both similar in magnitude across rating categories and evolves similarly over time across category. The equivalent coefficient for Financial Issuers ($\beta_{1,F}$ in Equation 2.1) follows a similar evolution over time (drops lower than Non-Financial parameter and recovers slower, seen in Figure 2.7) for AAA-, AA- and A-rated bonds. It is important to emphasize that the small sample size of AAA-rated Financial issuers (bonds) post-2010 is responsible for the volatility in the AAA-rated coefficient; for example, on 15-09-2011 only twelve such securities are present in the dataset. The BBB-rated coefficient ($\beta_{1,F}$) is very different from the other rating categories, before,

during and after the credit crunch.

2.6.2 Modelling Credit Spread

Markit's Annual Benchmark Spread is used as the measure of credit spread. Similar to the Bid-Ask Spread model, some model parameters (γ_k) are reviewed for several sub-models (by rating and date, as in Equation (2.2)). Referring back to Table 2.2 to aid in the interpretation, please note that the Relative Bid-Ask Spread is modelled as the exponential of the residuals of the regression equation. Figure 2.11 shows the γ_k coefficients for two indicator variables; Non-Financial issuer (Financial issuer) and Senior (Subordinate) issuer. Bonds issued by Financials and Non-Financials, *ceteris paribus*, are expected to trade at similar prices prior to the crisis.




Figure 2.11: Gamma coefficient for Non-Financials (left) and Senior (right) bond indicators.

Given the relative instability of the financial services industry (particularly banks) during the crisis, Financials are expected to trade at lower prices / higher spreads (yields) after the Northern Rock bank run (14-09-2007). Figure 2.11 shows that yields of Financial- and Non-Financial issuers, *ceteris paribus*, started to diverge at the time of the Northern Rock bank run and are yet to recover fully to pre-crisis levels. Similarly, Senior and Subordinate bonds traded at similar prices until mid-2007, but have since diverged. The negative coefficient for the relevant gamma coefficient,

implying lower yields for senior bonds, is not surprising, nor is the fact that the coefficient was close to zero; in a low default regime, recovery rates (affected by seniority status) are not likely to be an important determinant of bond prices. In a regime with high (perceived) default risk (premia), especially for Financial issuers, which issue most subordinate bonds, recovery rates are more likely to impact an investor's decision.




Figure 2.12: Gamma coefficient for Non-Financials across rating categories.

Looking at the Non-Financial indicator across rating categories (Figure 2.12), similar patterns emerge; the AA- and AAA-rated model coefficients follow the A-rated coefficients from Figure 2.11 very closely, where the AAA-rated model coefficient shows a high variance for the post-2010 period as a result of a small sample size. The BBB-rated model coefficients seem slightly different with coefficients far lower than the other rating categories pre-crisis, far more volatile estimates and an extreme reaction around the time of the Northern Rock event. Its coefficient remains below the coefficient of other rating categories for the entire period under study.




Figure 2.13: Gamma coefficient for log Duration (NF) on the left and RBAS on the right.

Figure 2.13 shows the evolution of the gamma parameters for Duration and the relative liquidity proxy, RBAS. The gamma coefficient of Duration (NF) is positive for most days during the observed time period, indicating a rising Credit Spread curve. Upon close inspection of Figure 2.13 (left), it becomes clear that the Duration parameter started its steep drop just days/weeks before the indicated Northern Rock Event (14-09-2007). Lastly, the zero/negative value of this parameter indicates a flat or falling credit spread curve which could be interpreted as the market trading on price rather than yield, where short term concerns over the value of investments dominate an investor's behaviour. Given the specification of Equation (2.2), the gamma parameter of RBAS (Figure 2.13, right) is directly related to the size of the liquidity premium that is extracted in the next section. At this point it suffices to observe that the relative liquidity proxy is positively priced on all days during the sample period, except for a very brief period during 2006/2007.




Figure 2.14: Gamma coefficient for Coupon (left) and Collateralized indicator (right).

Figure 2.14 shows the evolution of the gamma coefficients for Coupon (left) and the Collateralized indicator (right). The gamma parameter for Coupon rate is positive throughout the sample period, indicating that bonds with higher coupon rates trade at higher spreads. Please note coupon rates are expressed as whole numbers; e.g. the effect of a 5% paying bonds would be $5 \times \gamma$. This is also according to expectations and in line with literature (for example Leland, 1994) that speaks of a ‘tax-effect’, where the underlying idea is that bonds with a low coupon rate have a more favourable tax treatment than high coupon paying bonds. Collateralized bonds would also be expected to trade at lower credit spread, which is what is observed for most of the sample period, although with some variability (Figure 2.14, right). The zero/negative coefficient from 2004-2006 is unexpected at first sight but might be explained by the dynamics of supply and demand for, for example, mortgage backed securities.

Finally, the fit of the model from Equation 2.2 is considered by looking at the explained variation in credit spreads. Figure 2.15 shows the aggregated weekly R^2 -statistic over time for all four rating categories. The R^2 -statistic is a commonly used indicator of goodness-of-fit in linear regression (OLS) and is defined, most commonly, as the ratio of explained variance (variance of model’s predictions) to the total variance (sample variance of dependent variable).

$$R^2 \equiv 1 - \frac{SS_{reg}}{SS_{tot}},$$

where SS_{reg} and SS_{tot} are the regression and total sum of squares.

As can be seen, the model describes the data very well, but varies by both rating and time. Two additional observations can be made; values for R^2 are high in general, and between 2009-2013 the R^2 of A- and BBB-rated bonds seems to be substantially higher than the R^2 of the AAA- and AA-models.




Figure 2.15: Variation of the coefficient of determination, R^2 , for the credit spreads model over time (weekly) and by rating category.

The model coefficients presented in this section have focused on the A-rated bonds, but comparisons across rating have been made where most appropriate (Figure 2.12 and Figure 2.15, for instance).

2.6.3 Liquidity Premium Estimates

As remarked earlier, the liquidity premium will be investigated both in number of basis points and as a proportion of total credit spread. Whereas this section will focus on the numerical results for A-rated bonds in particular, similar results for other rating categories can be found in Appendix A (replications of Figure 2.16). For A-rated bonds, Figure 2.16 shows the following:

- **(left)** shows a time varying decomposition of the median credit spreads into a liquidity (black) and non-liquidity component (grey)
- **(middle)** shows the liquidity component in bps
- **(right)** shows the liquidity component as proportion of spread




Figure 2.16: Decomposition of credit spread (left) for A-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).

Since liquidity premia in Figure 2.16 are shown to vary over time, the liquidity premium of an A-rated bond of average liquidity is said to be time dependent. The time dependency of liquidity premia is not limited to basis points (if Figure 2.16 (right) were constant, liquidity premia would simply move proportionally with credit spreads), but extends to the proportion of credit spreads. In the pre-crisis period average liquidity premia appear low (relative to the rest of the sample period) and somewhat volatile. Just prior to the start of the credit crunch (2006-2007), average liquidity premia were near zero (Figure 2.16 (right)) on low credit spreads in general (Figure 2.16, left). The onset of the credit crunch caused the liquidity premium to rise from near-zero levels to approximately 50% of credit spreads.

The non-liquidity component, consisting of both the expected default losses and a credit risk premium, also increased dramatically in bps. This increase could be contributed to a number of factors. First, within the economic cycle and in the

context of the crisis, short-term expected default probabilities would have risen even if a bond's rating was unchanged. Arguably, though, this could only contribute marginally to the overall increase. Second, investors' levels of risk aversion due to higher (*perceived*) levels of uncertainty might have increased significantly during the crisis, pushing up risk premia. Third, there might have been increased uncertainty in what future default probabilities and recovery rates would be. This additional uncertainty attracts its own risk premium which would, therefore, have risen during the crisis.

Liquidity premia (Figure 2.16, right) were relatively high and stable for several years during/after the credit crunch (Figure 2.16, right), irrespective of levels of credit spreads (Figure 2.16, left) and appear to have started to decline at the start of 2013.

Rather than looking at the average (point-estimate) of the liquidity premia over time, Figure 2.17 investigates the distribution of liquidity premia by plotting various percentiles of the daily distributions (in basis points) over time, on a monthly basis. In the left-hand plot, for example, an 80% quantile of 200 on a given date means that 20% of bonds had a liquidity premium of more than 200 basis points on that date.




Figure 2.17: Time-varying nature of four weekly quantiles of daily liquidity premium distributions, in basis points (left) and in proportion of spread (right).

The distribution of liquidity premia is tight pre-crisis, but widens substantially during the 2008-2013 period, only recently becoming tighter again. The skew of the distribution (long upper tail) is a direct results of the skewed distribution of RBAS as $\exp(\epsilon_{BAS})$ from Equation (2.1). Lastly, estimates of average liquidity premia across rating category are compared in Figure 2.18.




Figure 2.18: Monthly estimates of median liquidity premia across rating category.

Taking monthly estimates to remove most of the very short term volatility of the time series to improve legibility, Figure 2.18 shows that before the crisis, the four categories behaved similarly, with the exception of the AAA-rated bonds, which saw far smaller liquidity premia. All rating categories display very low premia (0% - 10%) from mid-2006 to mid-2007 and shoot up as a response to the credit crunch (again, AAA-rated bonds are the exception). After the start of the credit crunch, the A- and BBB-rated bonds appear to behave differently; whereas AAA- and AA-rated bonds return towards pre-crisis levels (AAA-rated slightly elevated), bonds of lower credit ratings see far higher liquidity premia for a prolonged period of time, starting to return to pre-crisis levels in 2013. In general, bonds with a lower credit rating have higher liquidity premia (as proportion of spread).

The extensive descriptive analysis of the market in Section 2.4 demonstrates to some extent the changes that have happened over the period under study. Whereas

Figure 2.18 illustrates the liquidity premia as they are ‘observed’ in the market it may also be useful to consider the liquidity premium over time of a fictional, representative bond. To that end, Figure 2.19 shows the estimated liquidity premium of a typical bond over time, where the bond is A-rated, has a duration of five years, is senior, Non-Collateralized, paying a 2% annual coupon, two years old and with an average notional amount outstanding. The Financial status of the issuer is varied as well as the value for RBAS.

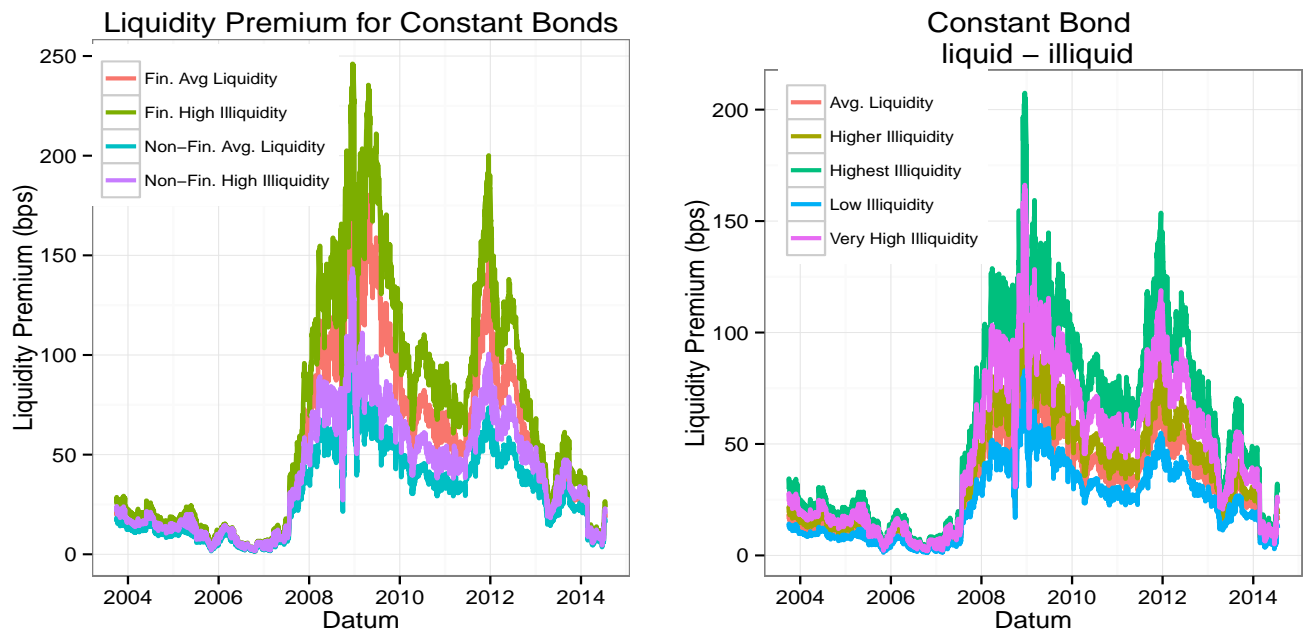


Figure 2.19: Predicted liquidity premia for constant bond over time (left), for various degrees of liquidity (right).

The general evolution of the graph in Figure 2.19 (left) corresponds to the observed liquidity premia reported earlier (Figure 2.16). The plot on the right illustrates to what extent, for the fictional bond, the liquidity premia would vary across various levels of liquidity; this plot corresponds to the observed liquidity premia quantiles from Figure 2.17.

2.7 Additional Analyses

Three additional analyses provide additional insight; the first is related to the choice and derivation of the new, relative liquidity proxy. The second additional analysis looks at the robustness of the model outcomes with respect to credit spread measure

and the section is concluded by investigating some of the properties of the RBAS.

2.7.1 Alternative Liquidity Proxy Specification

The Relative Bid-Ask Spread (RBAS) is probably, after an extensive literature review in Section 2.2, the only truly relative liquidity proxy, which exhibits the perhaps counter-intuitive property of having a constant average (expected) value, on a daily basis. Since RBAS is defined as the exponential of the residual term from Equation (2.1), its distribution, on a daily basis, is always expected to centre around 1 (standard log-normal for normally distributed residuals), irrespective of economic climate. The distribution for RBAS can either be wider or narrower, depending on the daily fit of the regression model (Equation (2.1)). Its relative nature and design does bring the attractive property of being uncorrelated with common bond characteristics. Using the same period and bond universe, a similar set of regression models is specified, but with the 'raw' bid-ask spread as liquidity proxy. The methodology to extract the liquidity premia is different and does not create a hypothetical perfectly liquid alternative; the method of premium extraction is based on and similar to the method used in Dick-Nielsen et al. (2012).

The use of bid-ask spreads, or indirect measures of the bid-ask spread such as the Roll-measure (as recently in Bao et al., 2011) or Imputed Roundtrip Costs (Feldhütter, 2012), have frequently been used to study the effect of illiquidity on asset prices. Figure 2.20 clearly shows that the time series of daily median bid-ask spread for investment grade bonds is highly time-, rating- and Financial/Non-Financial dependent, with the spread on financials increasing by much more during the crisis.




Figure 2.20: Shown on the same scale, the bid-ask spread for all IG ratings, both Financial and Non-Financial issuers, increased dramatically during the financial crisis.²

Again, regression models of credit spreads and bond characteristics are formulated (using identical covariates to Equation (2.2)), but including the bid-ask spread directly instead of RBAS;

$$\begin{aligned}
 (CS(i, r, t)) = & c(r, t) \\
 & + \theta_{1,FIN}(r, t) \times \log \text{Duration}(i, t) \times I_{FIN}(i) \\
 & + \theta_{1,NF}(r, t) \times \log \text{Duration}(i, t) \times I_{NF}(i) \\
 & + \theta_2(r, t) \times \log \text{Notional}(i, t) \\
 & + \theta_3(r, t) \times \text{Coupon}(i, t) \\
 & + \theta_4(\mathbf{r}, \mathbf{t}) \times \text{BAS}(\mathbf{i}, \mathbf{t}) \\
 & + \sum_k \theta_k(r, t) \times I_k(i, t) \\
 & + \epsilon_{CS}(i, t) \quad (\text{residual}).
 \end{aligned} \tag{2.4}$$

Since the dependent variable is credit spread rather than $\log(\text{CS})$, the coefficients and covariates can be interpreted directly to evaluate their contribution to credit spread in basis points. Instead of estimating the perfectly liquid equivalent bond, an

²Note that the abnormality for AAA-rated Non-Financials is due to a methodology change in the data; several Non-Financial issuers became Financial issuers on that day.

estimation procedure similar to Dick-Nielsen et al. (2012) is followed. The liquidity score for each bond is defined as $\theta_4(r, t) \times \text{Bid-Ask Spread}(i, t)$. Within each rating category (AAA, AA, A, BBB) and day, bonds are sorted on their liquidity score. Then, the size of the illiquidity contribution to the spread, for an average bond, is defined as the 50% quantile minus the 5% quantile of the liquidity score distribution in a particular bucket ($\theta \times (BAS_{50} - BAS_5)$). Therefore, the liquidity contribution measures the difference in credit spread between a bond of average liquidity and a bond that is very liquid. Compared to the approach of estimating perfectly liquid bonds, this measure is relative; the 5% quantile represents a very liquid bond on a particular day.

Comparing the time series of daily liquidity premium estimates for the A-rated bucket, with the estimates of median A-rated liquidity premia from the previous model shows (Figure 2.21) that the two move together somewhat, but are rather different. The alternative approach shows spreads of near zero for the entire pre-crisis period (bid-ask spread is not significantly priced), seems to react to the credit crunch more slowly, peaks around similar levels but drops much deeper in 2010-2011, increases drastically for only a few weeks during the European debt crisis from approx. 15% of spread to 65% of spread, only to return to near zero levels very quickly.




Figure 2.21: Comparison of A-rated liquidity estimates in basis points (left) and % of spread (right) for the proposed modelling approach (black) and the alternative approach (grey).

2.7.2 Robustness of Credit Spread Measure

The use of the Annual Benchmark Spread, in favour of other spread measures is described in Section 2.3 and on the one hand comes down to the tradability of the reference asset, but, more importantly can be attributed to reasons of backwards compatibility; other spread measures are only available after June 2009 in the Markit dataset. The asset swap margin and, in particular, the Z-spread and Option-Adjusted Spread are commonly used spread metrics that are only available after 2009. To show the extent to which the choice of spread measure influences the analysis in general and ultimately the estimated liquidity premia, a simple robustness check is performed by substituting in the new spread measures (using a post-2009 subset of the data) and documenting the results and potential differences.




Figure 2.22: Daily average values of three credit spreads measures for Non-Financials across four rating categories.

On aggregate, in Figure 2.22, spread measures appear to have very similar values and dynamics. Apart from some instability in the early days of publishing the Benchmark Curve and some discrepancies in the Z-spread after 2010, the spread measures appear very similar. Note that the Z-Spread and OAS are identical for bonds without optionality. Naturally, for modelling purposes, differences on the individual bond level are of most interest, rather than market-wide aggregates. Looking at individual bonds some discrepancies appear; in order to summarise how spread measures vary in one simple statistic, a straightforward metrics, namely a Mean Absolute Deviation, is used to summarize the data;

$$\text{MAD}_t = \sum_{i=1}^{i=N_t} \frac{|\text{Spread}_{s,i,t} - \text{BMS}_{i,t}|}{\text{BMS}_{i,t}} \times \frac{1}{N_t}$$

where s is the spread measure (either Annual Benchmark Curve, Z-Spread or OAS) for bond i on day t and N_t are the number of bonds in the universe at time t . The resulting series MAD_t is plotted in Figure 2.23.




Figure 2.23: Daily average values of three credit spreads measures for Non-Financials across four rating categories.

With a better understanding of the differences between spread measures for individual bonds and how these differences vary over time, the entire modelling process is repeated on a subset of the data (post 2009); this includes the modelling of the Credit Spreads, but not the modelling of Bid-Ask Spread in the first stage. Ultimately the interest of this Chapter lies with the different liquidity premia estimates. In Figure 2.24 liquidity premia are compared for each of the alternative spread measures; to aid interpretation, the liquidity premium using the Annual Benchmark Spread is used as a reference point.

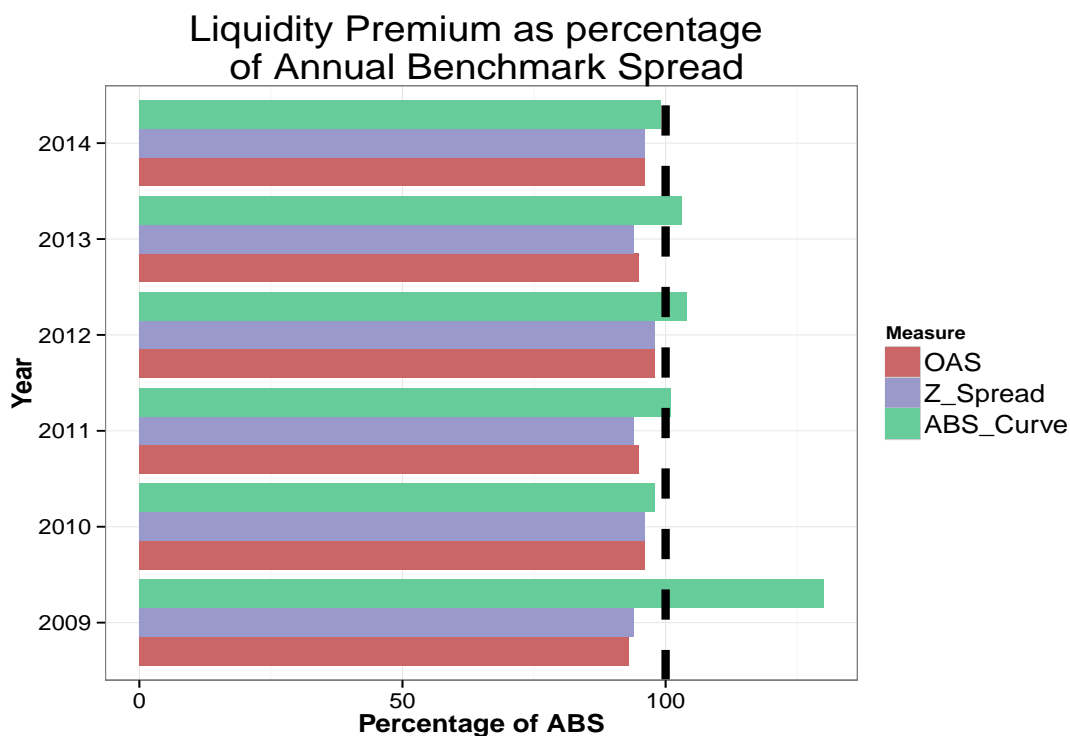


Figure 2.24: Comparison of liquidity premia estimates across various credit spread measures, aggregated by year to aid interpretation.

With the exception of the liquidity premium estimated in 2009 based on the Annual Benchmark Curve, for which very odd values are reported (see Figure 2.22), all spread measures lead to liquidity premia that are very close to the liquidity premia estimates using the Annual Benchmark Spread.

2.7.3 Investigating RBAS Properties

As remarked earlier, the liquidity proxy RBAS is entirely relative (daily distribution centred around one) and uncorrelated with common bond characteristics, all of which allows for the direct comparison of intrinsic bond liquidity. To gain a better insight into the properties of the relative liquidity measure, a set of Financial and Non-Financial issuers with multiple bonds outstanding on a particular day is used to graphically explore whether there seems to be evidence for an issuer specific liquidity effect. It is noteworthy that issuer specific liquidity has only been briefly explored by Dick-Nielsen et al. (2012), who considered an issuer specific liquidity proxy (non-zero trading days for issuer). Including additional covariates in the specified regression models, such as number of bonds outstanding by issuer or total notional

outstanding by issuers, yields a beta parameter that is largely insignificant and has been omitted from Equation (2.1). In Figures 2.25 and 2.26 the bid-ask spread (left) and RBAS (right) are shown for Financial and Non-Financial Issuers respectively.




Figure 2.25: Bid-ask spreads (left) and RBAS (right) for bonds issued by selected Financials. Ordering of bonds (by issuer and then by magnitude of BAS) on the left is preserved in the right-hand plot.

Two observations are important to make; first, daily BAS and RBAS for individual bonds/issuers are uncorrelated. The second observation is related to the issuer specific liquidity. Whereas issuer specific variables have been omitted from the model, Figure 2.25 appears to display a some issuer specific effect for Financials *during* the credit crisis. Issuer specific liquidity is defined as ‘generally more or less liquid than average’, where Figure 2.25 shows that issuers ABBEY, HSBC and LLOYDS seem to have most bonds outstanding with RBAS less than one. This effect is very limited for Non-Financial issuers (Figure 2.26), where perhaps the same issuer specific liquidity effect can only be observed for GE (General Electric).

The evolution of a bond’s relative liquidity over time is important to consider, as the distribution of the population of bonds does not change on a daily basis, it is of interest to find out how volatile the bond’s position is within that distribution. Estimates for RBAS are relatively robust over time, meaning that over short-medium periods of time, RBAS changes little. The volatility of RBAS is thereby mainly




Figure 2.26: Bid-ask spreads (left) and RBAS (right) for bonds issued by selected Non-Financials. Ordering of bonds (by issuer and then by magnitude of BAS) on the left is preserved in the right hand plot.

dependent on volatility of model parameters β_k in Equation 2.1, which are robust over short periods of time, and dependent on the movement of the bond's Bid-Ask Spread; both move in response to a changing market (model parameters) and idiosyncratic shocks. Figure 2.27 shows the evolution of weekly bid-ask spreads (left) and RBAS (right) of three bonds over a long period of time (multiple years) and it is clear that these bonds, despite short term volatility, operate at three different points of the RBAS spectrum. Please note the particular issue of the Royal Bank of Scotland (Figure 2.27 (third row)), which in recent years appears to be consistently more liquid compared to identical bonds which is likely the result of the government backing of RBS.

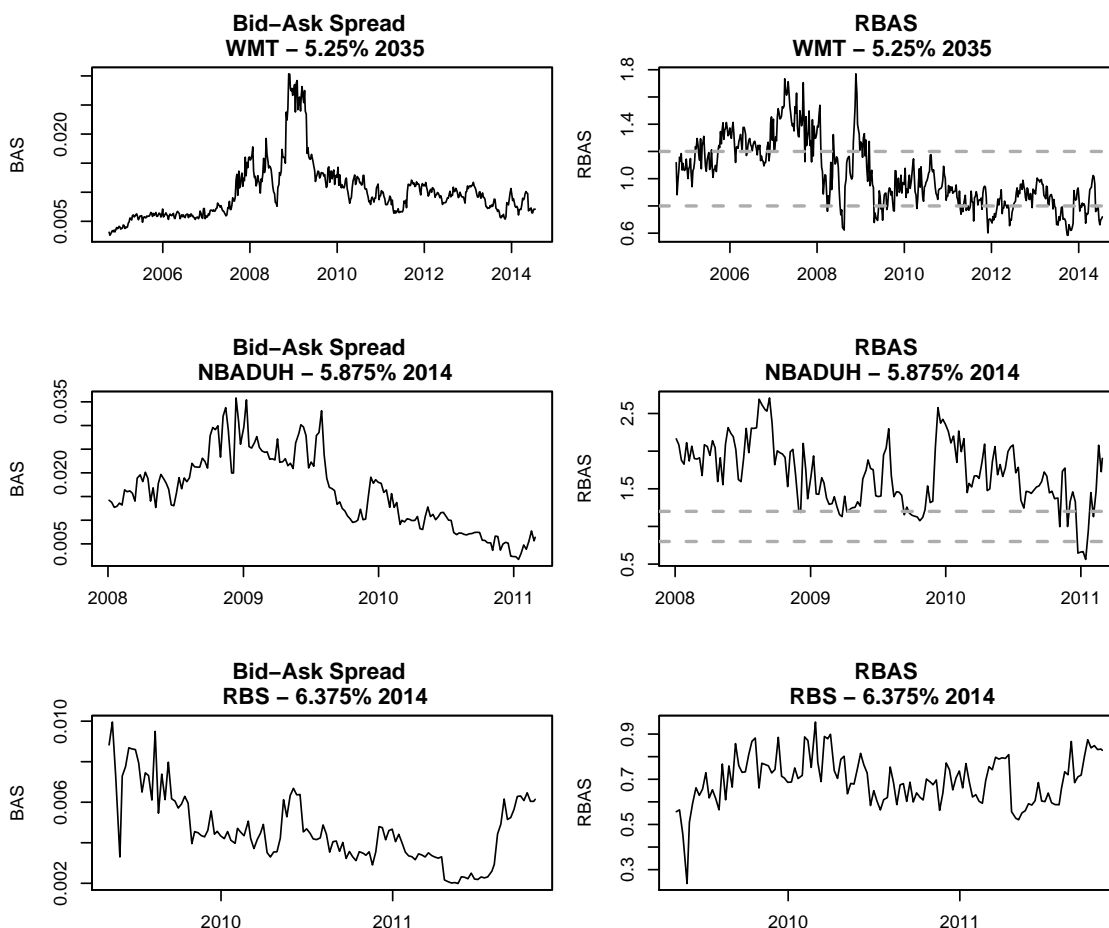


Figure 2.27: Bid-Ask Spreads (left) and RBAS (right) for three individual bonds over time.

2.8 Bank of England's Structural Model

The Bank of England published an article in their quarterly bulletin about decomposing credit spreads and liquidity premia (Webber, 2007), for which the details can be found in an earlier working paper by Churm and Panigirtzoglou (2005). The working paper discusses how the Leland and Toft (2006) structural model is used to arrive at historical default probabilities on the one hand, and is used to derive ‘contemporaneous forward-looking’ estimates of the fair credit spread on the other hand, which in turn dictates the estimated liquidity premium. This section aims to discuss some of the results from the Bank of England’s work and perform a sensitivity analysis on the structural model to illustrate the subjectivity and difficulty to arrive at decent parameters for the model. It also serves as an interesting example, given that Chapter 4 of this thesis is concerned with model uncertainty and

parameter risk.

2.8.1 Quick Model Overview

At the centre of the modelling effort in the working paper, is the Expected Default Frequency (EDF), directly taken from the Leland and Toft (1996) paper, with the EDF up to time t given by;

$$EDF_t = N\left[\frac{-b - (\mu - \delta - 0.5\sigma^2)t}{\sigma\sqrt{t}}\right] + e^{-\frac{2b(\mu - \delta - 0.5\sigma^2)}{\sigma^2}} N\left[\frac{-b + (\mu - \delta - 0.5\sigma^2)t}{\sigma\sqrt{t}}\right] \quad (2.5)$$

with a list of parameters,

- r , risk-free rate
- σ , asset volatility
- T , debt maturity
- α , bankruptcy costs
- P , leverage
- τ , tax advantage
- C , coupon yield
- δ , payout ratio

where $b = \ln(\frac{V_0}{V_B})$ and V_b , the default trigger, is described by;

$$V_b(C, T, \alpha, \delta, \sigma, \tau, P, r) = \frac{\frac{C}{r} \frac{A}{rT} - B - \frac{AP}{rT} - \frac{\tau Cx}{r}}{1 + ax - (1 - \alpha)B}.$$

The endogenous default boundary is computed using its subcomponents;

$$A = 2ae^{-rT} \mathcal{N}[a\sigma\sqrt{T}] - 2z \mathcal{N}[z\sigma\sqrt{T}] - \frac{2}{\sigma\sqrt{T}} n[z\sigma\sqrt{T}] + \frac{2e^{-rT}}{\sigma\sqrt{T}} n[a\sigma\sqrt{T}] + z - a$$

$$B = -(2z + \frac{2}{z\sigma^2 T}) \mathcal{N}[z\sigma\sqrt{T}] - \frac{2}{\sigma\sqrt{T}} n[z\sigma\sqrt{T}] + z - a + \frac{1}{z\sigma^2 T}$$

$$z = \frac{((a\sigma^2)^2 + 2r\sigma^2)^{\frac{1}{2}}}{\sigma^2}$$

$$a = \frac{r - \delta - (\sigma^2/2)}{\sigma^2}$$

$$x = z + a$$

In the Bank of England article, the above expression for the EDF is used to describe historical default data from the United States. The parameters of the model are chosen to best reflect a long running history. For instance, r is set to 0.08 (8%) as the average long-term (ten-year) government yield, Leland (1994) is followed in setting the bankruptcy costs to $\alpha = 30\%$, the average maturity set to five years and the leverage set to 41% as determined in a study by Standard and Poor's in 2001. The average historical credit spread for an investment grade bond is set to 136 basis points based on the Merrill Lynch Index.

Rather than producing EDFs, a calibrated model can also be used to generate credit spreads. The value of a coupon bond with semi-annual coupon payments K is given by the following equation:

$$B_T = \sum_{t=1}^T \frac{K}{2} e^{-rt} (1 - (1 - R)EDF_t(b, \mu, \delta, \sigma)) + \frac{1 + K}{2} e^{-rt} (1 - (1 - R)EDF_t(b, \mu, \delta, \sigma))$$

The par coupon spread is the value of K that makes the bond value equal to 1 (par value). This optimisation is solved using a Newton-Raphson method and the resulting value for K annualised. The credit component of the spread (both expected default and the risk premium) is computed by subtracting the risk-free rate; $K - r$. Using the above historical calibration, the fair credit spread comes to 83bps, which is 61% of the observed credit spread of 136bps.

To arrive at contemporaneous estimates of the credit spread, the Bank of England working paper uses several datasets and modelling choices to arrive at monthly estimates for several parameters of the model. The Merrill Lynch Global Index System is used to collect monthly credit spread estimates on UK corporate bonds, firm-level 1-year option implied volatility is collected from Bloomberg and a three stage dividend discount model is used to estimate the cost of equity (equity risk premium) using data from the Institutional Brokers Estimate System (IBES). The equity risk premium is based on firms with trading equity in the FTSE 100, which far from covers the entire Merrill Lynch bond index. The equity implied volatility does not have the correct maturity, not all firms in the index have options trading

on its equity and some do not have trading equity at all and need to be excluded (implied volatility data is available for 62% of the index constituents). To transform equity volatility to asset volatility a set of simultaneous equations from the Merton model is used to solve for σ numerically, but the Leland and Toft model is not equivalent to Merton, and this is only an approximation. The Bank of England paper discusses the difficulty in deciding whether an average or time-varying value for leverage is best to use; long-term average is used as target³. Recovery rates and bankruptcy costs are assumed to be fixed values and the long-term average leverage is derived from a sample of US equity. The choices for calibration are abundant and difficult to make due to the complexity of several parameters themselves (equity risk premium and asset volatility) or the inability to access reliable data at a granular level (earnings growth, leverage or bankruptcy costs).

Acknowledging the difficulty in deriving model parameters, the effect some of the parameters have on the estimated long-term investment grade liquidity premium (which translates directly to fair credit spread and in turn to liquidity premia) is substantial.

Apart from the model uncertainty when it comes to modelling something as complex as the liquidity premium, the structural models seem to be riddled with parameter risk, as there are many (unobservable) parameters to pin down. Added in the complexity of deriving some of the parameters and the difficulty to access reliable and frequently updated data, makes for a very complex model to parametrise; attempting to derive liquidity premia for individual issues rather than a deriving a market-wide estimate adds another layer of complexity.

³Long-term averages are used as those would serve well as leverage targets. The study fails however to distinguish between the leverage of financial and non-financial firms, which due to the business models is likely to differ irrespective of solvency risk. The leverage figures are based on US estimates.

Chapter 3

Quantitative Factor Investing in the UK Corporate Bond Market

3.1 Introduction

The modelling efforts presented in Chapter 2 fall under the category of statistical models that attempt to extract the liquidity portion of the credit spread. The models do so by building a new, relative liquidity proxy and estimate the credit spread of an artificially liquid bond. The estimated figures for the liquidity premia are conditional on a strict hold-to-maturity basis, which may not be feasible, or desired from a risk management perspective, yet is generally accepted among life insurers (annuity providers). This Chapter aims to introduce an alternative, nearly model-free approach to deriving liquidity premia using an intuitive interpretation of liquidity premia. Whereas the previous statistical model does not suffer from some of the obscure data and parametrisation requirements structural models carry, one could argue that it does describe a model in which a rather arbitrary and fictitious perfectly liquid alternative bond is created to derive premia. This somewhat theoretical interpretation and derivation of the liquidity premia might still benefit from an even simpler derivation based on a simpler interpretation of liquidity premia. If we consider liquidity premia to present itself as higher returns (as an accrued, expected, present value of credit spreads), liquidity premia can simply be observed by comparing the returns of relatively illiquid bonds and the returns of bonds that

are relatively liquid.

Again, the same Markit iBoxx dataset is used to compute liquidity premia in this Chapter, which makes full use of the dataset that is available by taking the opportunity to not only evaluate the risk and return characteristics of various definitions of an illiquidity factor, but extend factor investing principles to the wider dataset and aims to provide a review of quantitative factor investing in corporate bonds.

The topic of alternative/strategic/smart beta is often misrepresented in popular finance articles and may be abused as a marketing term rather than an accurate description of a group of investment types or strategies. According to Morningstar (Johnson et al., 2016), 950 such products are traded with an estimated value of more than 475 billion US dollar (December 2015), in an effort to cater to investors by promising alternative strategies with higher returns or cater to investors by allowing them to alter the relative risk profile of their portfolio.

‘Smart beta’ generally refers to passively following an index in which the index weights are not proportional to their market capitalizations, but based on alternative weighting rules. The alternative weighting can be designed to strategically expose the portfolio to factors. A factor can be thought of as any characteristic shared by or pertaining to a group of securities that is important in explaining their return and risk and is typically presented as a systematic risk premium, but can be related to anything including for instance, employee satisfaction (Blackrock, 2015). The Exchange Traded Fund (ETF) market has been instrumental in the rise of easily traded products, some of which are classed “smart beta”. Some of the factors that funds attempt to capture might be difficult to characterise as “smart” products, as most are geared towards capturing factor risk premia that are well-known. Investors need to be aware of smart product ETFs that try to justify higher management fees based on overly complex strategies and the ambiguous position of alternative beta products as both an active and passive investment (strategy). Smart products are active in design, where the ‘intelligence’ of an alternative index is created and is passive in its implementation as no day to day decisions about portfolio positions is

necessary and mostly quantitatively, pre-defined criteria rebalance portfolios.

As alternative beta funds generate tremendous inflows and subsequent (media) attention, these new ETFs have also received their fair share of criticism. The criticism ranges from sceptics regarding the validity of factor outperformance versus higher implied risk, over-specification or data dredging to create strategies without longevity, and the reasoning that the market anomalies, if behavioural biases, that strategic funds systematically might attempt to exploit will simply cease to exist due to incredible inflows, leading to an almost self-fulfilling prophecy. With cheap index funds at anyone's disposal and the realisation that naive allocations tend to outperform more sophisticated allocations over long periods (DeMiguel et al., 2009), the justification of high fees continues to be challenged.

A large body of academic research highlights the empirical finding that long term equity performance can largely be explained by factors; for example, five well-documented factors associated with the equity risk premium include Value (Basu, 1977; Reinganum, 1981), Low Size (Banz, 1981), Low Volatility (Jensen et al., 1972, Haugen and Heins, 1975), Quality (Sloan, 1996) and Momentum (Jegadeesh and Titman, 1993; Carhart, 1997). Long term equity performance can be captured by these systematic risk premia, which alternative beta funds attempt to do, where the possibilities for the operationalisation of these risk premia are endless.

The remainder of this Chapter is organized as follows: Section 3.2 will briefly revisit the Markit iBoxx GBP Corporate Bond dataset; Section 3.3 discusses the portfolio methodology (of the replication) as well as the practical implementation of the factors, Section 3.4 discusses the distributional properties, (risk-adjusted) performance, turnover and transaction costs of single-factor portfolios, where Section 3.4.5 is specifically dedicated to a discussion about liquidity premia. Finally, Section 3.5 studies multi-factor portfolio construction and optimisation under several strategies.

3.1.1 Review of Credit Specific Literature

The body of literature for factors in equity risk premia is extensive and spans several decades, with the most well-known work by Fama and French (1993) who extend a standard Capital Asset Pricing Model (CAPM) to include factors beyond the market factor. Whereas the empirical evidence in equity markets is extensive, the research on corporate bond factors is far more limited. The first work to jointly analyse multiple factors using a unified approach was by Houweling and Van Zundert (2014). In addition to this work, there are also some contributions on Low-Risk (Quality), studied by Frazzini and Pedersen (2014), and Value, studied by Correia et al. (2012).

Correia et al. (2012) and Frazzini and Pedersen (2014) evaluate the existence of individual factor outperformance. Specifically, Correia et al. (2012) evaluate the predictive power of accounting-based and equity market based information in relation to corporate credit spreads. Incorporating these inputs leads to a structural model that is best able to forecast default, and, interestingly, the market does not completely price this and does so with a lag; Correia et al. (2012) capture the possibility to take advantage of value-investing to achieve excess returns. Frazzini and Pedersen (2014) construct a ‘BAB’ (Betting-against-Beta) factor which is based on the driving factor of leverage constraints in their model for US equities, 19 international equity markets and other asset classes, including US-Treasury bonds and corporate bonds. Across all asset classes they find that a self-financing, market neutral portfolio, long in low-beta assets (levered to one) and short in high-beta assets (de-levered to one) produces significantly higher Sharpe ratios and positive alpha. The existence of Momentum in corporate bonds has recently been documented by Jostova et al. (2013), who find that non-investment grade bonds carry Momentum profits of 192bps, with bonds placed by private firms commanding even higher Momentum profits of 282bps. They fail to find evidence of momentum profits for investment grade bonds, but do find that while some momentum spills over from equities to both investment and non-investment grade bonds, there are large and significant bond-specific momentum effects in non-investment grade bonds. The treatment in Houweling and Van Zundert (2014) covers multiple factors studied in

the US corporate bond market, in both the investment grade and high-yield segment. They find statistically significant premia for Size and Value factors in both market segments; out-performances for Size being 75bps and 402bps for investment grade and high yield respectively. Their ‘Low Risk’ factor does not outperform the general market on a strict return basis, but due to its low levels of volatility, results in a significantly higher Sharpe ratios. Houweling and Van Zundert (2014) continue to consider the value of corporate bond factors as part of a diversified portfolio of asset classes and show positive alpha of the investigates corporate bond factors exists.

This Chapter builds on the relatively small body of literature in several ways. Using a similar approach to Houweling and Van Zundert (2014) in identifying and constructing the factor portfolios, seven factors, rather than five are considered, adding High Volatility and Illiquidity to the factors considered in Houweling and Van Zundert (2014). The research presented in this Chapter is the first comprehensive study of seven factors using a common methodology applied to the same UK dataset, in which the risk and return characteristics and relative risks of each factor portfolio are reviewed.

The Illiquidity factor is investigated in greater detail and used in interpreting liquidity premia for corporate bonds. Not only are the risk and return characteristics for the factor portfolio discussed, liquidity premia are also derived for subsets of the market, where particular interest in this application lies in the risk characteristics of the returns associated with illiquid assets. The relative merits of the proposed modelling methodology are discussed and both the size of the premia as well as its dynamics over time are of interest.

Existing literature, in both equities and credit generally employ a fixed holding period for portfolios when working with quantitative factor portfolios, whereas the proposed methodology allows for changes to the portfolio to be made on a monthly basis. Rather than fixing the holding period of portfolios for a typical six or twelve months, i.e. yearly rebalancing according to some set of quantitative criteria, and rather than immediately trading bonds that fail to meet the specified factor criteria,

a tolerance level is specified, effectively controlling the portfolio's turnover at the expense of diluting the factor's exposure. An optimal tolerance level for each of the factor portfolios is studied, maximising return net of transaction costs. Rebalancing a corporate bond portfolio at random times using a tolerance parameter rather than at fixed intervals adds a degree of realism to the proposed methodology as fund managers would typically like the opportunity to re-balance at frequent intervals.

The UK corporate bond market, even the investment grade segment, is an illiquid market resulting in significant trading costs. Not only the extent to which transaction costs, explicitly defined as bid-ask spreads, eat into the strategy's return (investigating post-transaction levels of alpha) is investigated, but also how the tolerance level can be set to achieve an optimal holding strategy. This represents a trade-off between factor dilution (at high tolerances) and portfolio turnover / transaction costs. Evidence of factor dilution exists for most factors and it is observed that the suggested, optimal holding strategy allows for a substantial tolerance level, realising limited portfolio turnover.

Motivated by a relatively high degree of benchmark risk for single-factor portfolios, that is, the risk associated with the outperformance versus the market, the benefits of allocating to multiple factors are studied in some detail. It is paramount that a strategic asset or factor allocation should first and foremost be driven by an institution's objectives and constraints, yet three strategies and several implementations are formulated under which multi-factor allocations might be used. Each of the strategies represents a general, potential objective for pursuing a multi-factor corporate bond allocation: enhance risk-adjusted returns, limit downside risk and limit relative risk. Multi-factor allocations allow higher returns, lower risk, reduced cyclicity and lower relative risk through diversification of outperformance of single factor portfolios.

The period under study is relatively short (11 years) and includes several years of severe financial distress. Since factor investing is concerned with excess returns that materialise over longer periods of time, and market participants in the corporate bond market can have investment horizons that span several decades, some cursory

insight into the variability of optimal portfolio weighting schemes is provided.

3.2 Description of the Data

Using the same Markit iBoxx dataset for Investment Grade corporate bonds in the UK, discussed in detail in Chapter 2, granular data is available, updated with high frequency and supplemented with many bond characteristics at any given time. Unfortunately, the dataset is not free from survivor-ship bias as bonds can leave the dataset prior to maturity without accurate information about what happened. To account for bonds being downgraded to High Yield and bonds defaulting, the following heuristics are used to compute a final return value on the day the bond disappeared. As the dataset provides no information about what happened to the bond, it is assumed to either have defaulted, downgraded to high yield status or become unrated. The probability of each of these outcomes depends on the rating category prior to disappearance, and is represented by a credit rating migration matrix (see Table B.1 in Appendix B). To compute the final return, a probability weighted average (based on Table B.1) of losses associated with each event is used. For the default and NR event a fixed recovery rate of 35% of the last observed price is assumed, for downgrades to High Yield a loss of 25% is assumed, irrespective of initial rating or time to maturity. It is important to note that while these adjustments are extreme simplifications, the impact of this exercise is rather small. Figure B.1 in Appendix B contains the number of bonds that disappear from the dataset for each of the factor portfolios.

When evaluating the performance of factor portfolios of corporate bonds, the analysis focuses on the excess returns over maturity-matched government bonds. Each bond is assigned a benchmark-Treasury bond by Markit, matching the time-to-maturity of the corporate bond, and the excess return of the corporate bond is provided as one of the analytical values in the dataset.

The daily excess (simple) return is provided in the dataset. This accounts for cash payments (coupons) and is based on quoted bid prices (Markit, 2014). These daily excess simple returns are converted into log returns and aggregated to monthly

returns. Using excess returns effectively removes the Term risk premium, which can be captured by investing in Treasuries alone; the emphasis lies in capturing the default/transition premium or credit (risk) premium element of a corporate bond's return.

3.3 Methodology & Risk Factors

The general strategy employed by the factor portfolios is to invest in a fraction of the market with particular attributes; the fraction of the market with high 'factor exposures'. The Value factor for example, would invest in bonds that are considered undervalued according to some pre-specified, quantitatively defined, criteria. This is an alternative weighting scheme that moves away from the traditional market portfolio. In Section 3.3.1 the quantitative criteria that determine a bond's factor exposure at any given time are discussed. The initial portfolio consists of those bonds that score above the 90th percentile on the factor exposure. This uses the top 10% exposure, which follows Houweling and Van Zundert (2014) and is a common definition in equity market factor modelling. In the famous extensions of the CAPM by Fama and French, (1993), the factor exposures are modelled by going long equities with high factor exposures and shorting the equities with the lowest factor exposures. Their additional factors (in the three-factor model) are aptly named 'Small-minus-Big' and 'High-minus-Low', referring to the long/short positions taken in small and large caps, and equities with high and low price-to-book values, respectively. Whereas shorting large cap equities on a major stock exchange sees no constraints, the shorting of corporate bonds is very expensive and cumbersome, and often no market will exist whatsoever. In this respect, the factor construction therefore deviates from the common practice by Fama and French (1993).

The initial portfolio, containing the top 10% bonds with the highest factor exposure, is then rolled forward in time. In the case of the UK investment grade market this translates to a portfolio of around 140 different issues, at any time. Whereas previous literature has taken a fixed holding period approach (albeit with overlapping portfolios (Jegadeesh and Titman, 1993), a portfolio approach is considered

in which changes to the (factor) portfolio can be made on a monthly basis. This strikes a balance between maintaining factor exposure, i.e. trading a bond when it fails to meet the specified criteria for inclusion and a fixed re-balancing scheme that is entirely inflexible. This ‘semi-continuous’ approach relies on the factor specification (10% of bonds with highest exposure, market-value weighted) and a tolerance level, denoted by λ . Therefore, rather than switching all bonds that are no longer in the top 10% of a factor’s exposure in a subsequent month, the tolerance controls how strictly defined the factor remains. Only if a bond falls outside the top 10% + λ , is a bond replaced by the bond with the highest exposure to the factor among those that are currently not in the portfolio. Higher levels of λ will prevent excessive turnover, but might also lead to diluted factor exposure, an interesting optimisation hypothesis that is investigated in the remainder of this Chapter.

3.3.1 Risk Premia Factors

The risk factors described in this section represent well-established concepts in equity markets and are observed to exist in many international equity markets. Literature describing how these premia become to exist from a, usually, behavioural, supply/demand or systematic risk viewpoint is also well-established and is briefly mentioned here for each of the factors. The exact operationalisation of the factors is subject to debate, naturally; whereas one, rather arbitrary, method of using a proxy for the factor is chosen, this can be done in numerous ways. However, at least in equity markets, there is strong evidence that over longer time horizons, these factors, despite their simplistic application here, have led to abnormal returns. The well-known concepts introduced here, many part of standard finance textbooks, are therefore by no means implemented as ‘smart beta’, yet have a proven track records as constructs (in equity markets), spanning many decades.

Size The small cap premium on equities is perhaps the most widely acknowledged factor commanding a risk premium and has been shown to exist in both developed and emerging markets. The source of the premium in equity markets can be seen as either an increased systematic risk under the efficient market hypothesis or as a

proxy for other underlying risk factors such as illiquidity (Amihud, 2002), default risk (Vassalou and Xing, 2004) and information asymmetry (Zhang, 2006). Size is typically defined as the index weight of a company in a given month; the market value of the issuer's outstanding debt as a proportion of the total value of all outstanding debt is used in this study.

Following Houweling and Van Zundert (2014), this reflects the size of a company's public debt and is directly related to the often-used definition of the Size factor in equity investing, where total market capitalization can be seen as a proxy for incomplete information (Daniel and Titman, 1997; Van Dijk, 2011). An alternative definition of the factor could be, for instance, the market capitalization of the firm's (publicly trading) common equity. The Size factor invests in companies with the smallest weight in the index.

Value Value investing is a well-documented concept in equity markets and is, in essence, concerned with buying company stock that is priced cheaply relative to the fundamental value of the firm, using some criteria. The Value factor in the Fama-French framework (Fama and French, 1993) is referred to as the High-Minus-Low (HML) factor and represents the difference in performance between companies with high and low book-to-market values. Many alternative criteria, often based on company fundamentals, have been used to identify companies offering 'good value'. Correia et al. (2012) extend this principle to the bond market where they compare the company's implied riskiness (credit spread) to several measures of fundamental value, including distance-to-default (Merton, 1974), leverage and profitability. Since factors in this study are designed using bond characteristics alone, only the analytical values available in the dataset are used to identify bonds that appear to be relatively cheap. More specifically, the observed credit spreads are regressed on several risk drivers, similar to the modelling process in Chapter 2;

$$\log(CS_{i,t}) = c + \sum_{r=1}^4 \beta_r Rating_{i,r,t} + \beta_5 \log(Dur_{i,t}) + \beta_6 Sen_{i,t} + \beta_7 Col_{i,t} + \epsilon_{i,t}, \quad i = 1, \dots, N \quad (3.1)$$

where N is the number of unique bonds in the dataset, $Rating_{i,r,t}$ is an indicator variable equal to 1 if the bond's rating is r in month t , Dur_i is the bond's modified duration, Sen_i is an indicator variable (Senior / Subordinate) and Col_i is an indicator variable (Collateralized / Not-Collateralized). The cross-sectional regression in Equation 3.1 is estimated for each month (t) using standard OLS.

The bonds eligible to enter the Value factor are the bonds with the largest positive percentage deviation between fitted and observed (log) Credit Spreads in a given month; i.e., the bonds that appear to be undervalued according to the pricing equation, indicated by negative values of the error term in the regression equation ($\epsilon_{i,t}$).

Quality Quality investing is widely accepted as part of a 'fundamental analysis' or 'stock selection' strategy, but even in equity markets it is a relatively recent concept when defined in quantitative terms. Beside accruals, which was used as a proxy for quality by Sloan (1996), proxies in equity markets include earning persistence (Dechow et al., 2010), stable growth and a high payout ratio (Asness et al., 2013).

In line with the estimate of fundamental value for the previous factor and following the established literature that illustrates how bonds of short maturity (de Carvalho et al., 2014) and high rating (Frazzini and Pedersen, 2014) earn higher risk-adjusted returns than the benchmark, an analogous rating/maturity approach to constructing a Quality factor is deployed.

Similar to the approach taken in Houweling and Van Zundert (2014), all A- and BBB-rated bonds are excluded, and, each month, select M bonds (AAA- or AA-rated) with the lowest term-to-maturity that make up 10% of the total index. The resulting boundary value M_t varies over time (Figure 3.1).




Figure 3.1: Boundary value M_t indicating the highest time-to-maturity included in the Quality factor on each day. Tick marks indicate 1st January.

Momentum Momentum represents the excess return of securities with a stronger past performance. Over periods of time, securities might follow a trend whereby strong performers will continue to perform strongly. The evidence from equity markets comes from, for instance, Jegadeesh and Titman (1993) who observed that buying strong performers and selling weak performers in the US stock market during 1965 - 1995 produced abnormal returns. Carhart (1997) extended the three factor Fama-French model to include a momentum factor. Empirical research suggests the momentum effect is strongest over the next 3-12 months, which typically leads to momentum strategies having high turnover. The reason for the existence of a momentum effect is widely discussed and debated. The most cited theories are all behavioural; investors irrationally over- or under-react to news, the reactions driven by self-attribution, conservatism bias, overconfidence or aversion to realize losses (Barberis et al., 1998; Hong et al., 2000).

The Momentum factor has previously been explored for corporate bonds, but the results are far from conclusive. The results seem particularly divided when it comes to the distinction between Investment Grade and High Yield bonds. Khang

and King (2004) and Gebhardt et al. (2005) find evidence for a reversal effect for Investment Grade bonds that may be due to institutional herding (Cai et al., 2012) and Jostova et al. (2013) fail to find a momentum effect for Investment Grade bonds. In the High Yield segment of the market, the Momentum effect has been shown to exist and lead to excess returns (Jostova et al., 2013).

In this study Momentum is defined as the cumulative return over the past six months;

$$\text{MOM}_{i,t} = \sum_{s=t-5}^t R_{i,s},$$

where $R_{i,t}$ is the log return of bond i at month (time) t .

High & Low Volatility The Low Volatility factor seems at odds with the fundamental principle that higher systematic risk is associated with higher returns in a CAPM-world, where riskier assets (characterised by high beta) command higher required returns. The outperformance of low volatility stocks is well-documented and spans a range of volatility measures and international stock markets (see for example Clarke et al., 2006). Many behavioural explanations have been offered for the outperformance of low volatility stocks: representativeness, overconfidence, agency issue or asymmetric behaviour in bull/bear markets. Please refer to Sefton et al. (2011) for a comprehensive review for the behavioural biases that may underpin the low volatility phenomenon.

The equity market literature has documented the low volatility effect using a wide range of volatility estimates and is defined, in this study, as the rolling six month standard deviation of returns, as a proxy for volatility;

$$\text{VOL}_{i,t} = \sqrt{\frac{\sum_{k=0}^5 (R_{i,t-k} - \bar{R}_i)^2}{6}},$$

where \bar{R}_i is the average return over the six-month period ($\frac{\sum_{k=0}^5 (R_{i,t-k})}{6}$).

In addition to the Low Volatility factor, a High Volatility factor is explicitly defined, which is constructed in an identical, but opposite manner. The rationale for including the High Volatility factor is to challenge the rejection of a CAPM-type

world where risk (volatility) is not rewarded. It is expected that high volatility bonds would have suffered from great losses during the financial crises, but would have recovered from the crisis more rapidly. Lastly, the difficulty in trading high volatility bonds during periods of financial distress (2008-2009) is acknowledged, as market liquidity can disappear and the practical obstructions to trading these bonds are not to be underestimated.

Illiquidity The effect of illiquidity and (excess) stock returns on equities has been well-documented (Amihud, 2002; Brennan and Subrahmanyam, 1996). A related strand of literature has investigated the role of liquidity in the relatively illiquid market for corporate bonds (Bao et al., 2011; Chen et al., 2007). Amihud et al. (2015) investigate the addition of a liquidity factor IML (illiquid-minus-liquid) to the commonly used Fama-French factor framework and observe statistically significant outperformance across global economies. In this Chapter, an explicit illiquidity factor for corporate bonds is introduced. Using the results obtained in Chapter 2, bonds are sorted according to their relative liquidity. The newly derived relative liquidity proxy from the previous Chapter, the RBAS (Relative Bid Ask Spread), defined as the exponent of the residual of the regression equation (2.1) in Section 2.5.2 in the previous Chapter, is used directly. The detailed modelling process to arrive at a (relative) liquidity proxy is described in Section 2.5 in the previous Chapter.

Alternative definitions of an illiquidity factor could be equity market cap, trading volume, age of the bond or bid-ask spreads directly; all of which have been used as proxies for illiquidity directly (Houweling et al., 2005). The relative liquidity measure defined in van Loon et al. (2015) and Chapter 2 of this thesis benefits from being uncorrelated to common bond characteristics by definition; this means the ‘illiquidity factor’ is less likely to capture the underlying effects of, for instance, long-duration bonds.

In addition to being part of the wider quantitative factor analysis investigated in this Chapter, special attention is reserved for the illiquidity factor as a way of providing competing estimates of a liquidity premium. The illiquidity factor is

defined using the RBAS statistic introduced in Chapter 2, and is computed on subsets of the bond universe, in addition to the market-wide illiquidity factor.

3.4 Single-Factor Portfolios

In order to evaluate the extent to which factor portfolios historically earned a premium beyond the market-wide credit risk premium, a portfolio strategy with a tolerance (λ) of 40% is used, meaning that a bond will only be switched if it falls outside a wider definition of the factor criterion (outside top 50% exposure rather than top 10% factor exposure, as per the strict factor definition). The following sections discuss the non-normality of factor returns, risk-adjusted (out)performance, turnover, transaction costs and optimal choice of factor tolerance.

3.4.1 Non-Normality of Factor Returns

Before evaluating the (risk-adjusted) performance of the factors, the distributional properties of the factor and market portfolios are investigated. For a full review of risk, not only are several different measures of risk-adjustment specified, but most importantly, deviations from normality in the distribution which greatly affect tail risk, are specified. Ultimately, when constructing multi-factor portfolios, the traditional mean-variance approach is flawed when the return series are skewed and leptokurtic, properties of monthly financial returns that are, to varying degrees, present in all asset classes (see for example Jondeau et al., 2007; Kat and Brooks, 2001 and Rachev et al., 2005). Formally, the excess return series of each factor portfolio is subjected to a Shapiro-Wilk test of normality (Shapiro and Wilk, 1965) and extremely strong evidence against normality of all factor returns is observed¹. A skewed Student t -distribution from Fernández and Steel (1998) is fitted to the return series, in an attempt to better capture fat tails and skewness. Note that many 'skewed t -distributions' have been formulated and a review is beyond the scope of this Chapter; Aas and Haff (2006), Azzalini and Capitanio (2003) and Hansen (1994)

¹The p-value for the observed (excess) return series (N = 110 months); Market = 1.0e-05, Liquidity = 9.9e-06, Low Vol. = 7.3e-11, High Vol. = 2.1e-07, Momentum = 4.2e-09, Value = 0.0003, Quality = 1.9e-08, Size = 2.8e-06

provide a non-exhaustive overview of skewed t -distributions. Though the skewed t -distributions will better capture some of the non-normality of the observed return series, the choice for this particular (implementation) of a distribution somewhat arbitrary, though in line with common practice. Formal tests do not reject the skewed t -distribution for any return series ($p \geq 0.1$), which does not imply better or more appropriate distributions might exist. Figure 3.2 shows histograms of the factor distributions with fitted empirical Gaussian (red) and skewed t -distribution (blue) densities.

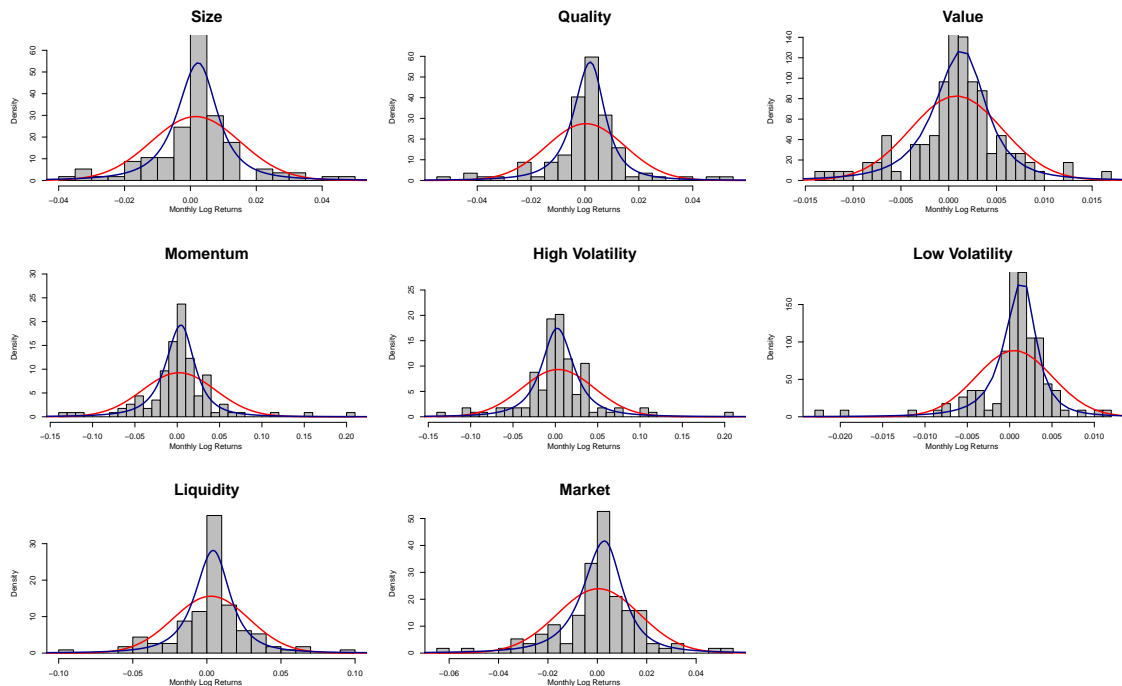


Figure 3.2: Distribution of returns for each of the factor portfolios and the market portfolio. The blue and red line correspond to the fitted densities of a skewed t - and Gaussian distribution respectively.

The estimated densities for the fitted Gaussian and fitted skewed- t differ considerably for all factors. Looking at the market factor, considerable kurtosis and a slight negative skew is observed. Naturally, the variance of the High Volatility factor is several multiples of the variance of the Low Volatility. Every factor portfolio, including the market portfolio, suffers from excess kurtosis.

For the purposes of risk management, the tails, and in particular the left tail of a distribution, are vital as this governs probabilities of large losses. The fat tails for the return distributions of the factor portfolios cannot be ignored.

3.4.2 (Risk-adjusted) Measures of (Out-)performance

The distributional properties of the factor returns suggest the need for appropriate measures to evaluate the risk-return trade-off for each factor. Before presenting the risk-adjusted performance of the factors in Table 3.1, the (risk-adjusted) measures of performance are briefly reviewed.

Risk is portrayed in numerous ways, and since each risk adjustment has its own merits and flaws, the resulting spectrum of risk measures considers riskiness in a most comprehensive way. Firstly, the Sharpe ratio, the most common risk adjustment, is used, where the Sharpe Ratio is defined as the expected return of the asset over the risk-free rate per unit of total risk, defined as the standard deviation. Since only returns in excess of Gilts are used throughout this Chapter, this reduces to;

$$SR_p = \frac{E(R_p)}{\sqrt{\text{var}(R_p)}},$$

where $E(R_p)$ is the expected return of the portfolio and $\sqrt{\text{var}(R_p)}$ its standard deviation. Next, the Omega ratio (Keating and Shadwick, 2002) is considered, defined as the probability weighted ratio of gains versus losses for a return target equal to the mean return of the market portfolio;

$$\Omega = \frac{\int_{E(R_p)}^{\infty} (1 - F_{R_p}(x)) dx}{\int_{-\infty}^{E(R_p)} F_{R_p}(x) dx},$$

where $F_{R_p}(x)$ is the cumulative distribution function of the factor portfolio. Investors would prefer assets, strategies or factors with high values for Ω , indicating a large proportion of gains relative to losses for threshold $E(R_p)$.

Value-at-Risk measures are considered by computing the VaR and the Conditional VaR (or TVaR, Expected Shortfall) at the 99th percentile, where TVAR is computed as;

$$\text{TVaR}_{0.99}(R_p) = E[R_p | R_p \leq -\text{VaR}_{0.99}(R_p)],$$

for a continuous random variable R_p .

It is important to note that the VaR and TVaR are computed using the fitted distributions (Gaussian and skewed-t) rather than historical return values. Given that some of the return distributions appear to exhibit non-normality (Section 3.4.1), the values for the Gaussian- and skewed-t based VaR measures can be very different; this reflects the tail risk present in most factors, but also the market factor.

Maximum Drawdown (MDD) is defined as the peak-to-trough decline of an investment during a specific period and is usually quoted as a percentage of the peak value;

$$MDD = \frac{P - L}{P}$$

where P is the peak value before largest drop and L is the lowest value before a new high is established.

Next, risk adjustments based on systematic risk are considered using both a CAPM framework as well as a Fama and French (Fama and French, 1993) factor model. The return of the factor portfolio (R_p) are regressed on the market portfolio;

$$R_{p,t} = \alpha_p + \beta_p R_{m,t} + \epsilon_{p,t}$$

where $R_{p,t}$ and $R_{m,t}$ are the return of factor portfolio p and the market portfolio in month t , respectively.

The value of the intercepts, α_p , are interpreted as a measure of a risk-adjusted outperformance, adjusting for the risk inherent in its exposure to the market (measured by β_p). The Fama-French factor model extends the 1-factor market model with four additional factors;

$$R_{p,t} = \alpha_p + \beta_1 \text{RMRF}_t + \beta_2 \text{SMB}_t + \beta_3 \text{HML}_t + \beta_4 \text{WML}_t + \beta_5 R_{m,t} + \epsilon_{p,t},$$

where RMRF_t is the equity market premium, SMB_t is the equity size premium, HML_t is the equity value premium, WML_t is the equity momentum premium and

$R_{m,t}$ is defined as before. Again, values of α_p are interpreted in a similar fashion, where positive α_p indicates an outperformance of the asset after accounting for its exposure to a set of systematic (equity) risk factors.

Measures of relative risk, the risk associated with outperformance, are also reported. The realized tracking error, defined as the standard deviation of active returns, provides insight into the volatility of the outperformance;

$$\text{TE}_p = \sqrt{\frac{1}{T} \sum_{t=1}^T (R_{p,t} - R_{m,t} - \bar{R}_p - \bar{R}_m)^2}$$

where $R_{p,t} - R_{m,t}$ is the active return of factor portfolio p at time t , relative to the returns of the market ($R_{m,t}$) at time t . Lastly, the information ratio is computed, for which higher values indicate a greater ability to out-perform given the relative risk assumed in the portfolio;

$$\text{IR}_p = \frac{\bar{R}_p - \bar{R}_m}{\text{TE}_p}.$$

3.4.3 (Risk-adjusted) (Out-) Performance of Factor Portfolios

Before moving on to the risk and return characteristics of the factor portfolios, it is useful to investigate the return series graphically. Figure 3.3 (left) shows the total return of each factor portfolio and the market portfolio; Figure 3.3 (right) shows the total return series of interest in this Chapter; the return over duration-matched Gilts. The duration matching is performed on an individual bond basis where each corporate bond that could be included in a factor portfolio is assigned (on each trading day) a matching Gilt. The ‘reference’ portfolio of duration-matched Gilts therefore varies for each factor portfolio. On an absolute basis, investing in the market portfolio would have led to a greater terminal value than investing in the Quality, Low Volatility or Value portfolio, but on a Term Premium adjusted basis (Figure 3.3, right) only, marginally, outperforms the Quality factor.




Figure 3.3: Total return for each of the factor portfolios and the market portfolio (left) and the total return in excess of duration-matched Gilts (right). Tick marks indicate the start of the year.

As can be reasonably expected, the factor portfolios are affected differently by the financial crisis in 2008/2009, which can be observed directly from the return series. Whereas the market portfolio dropped around 25% relative to duration-matched Gilts, the Momentum portfolio lost close to 50% and the Low Volatility market lost less than 10%. A summary of risk and performance measures can be found in Table 3.1; please note all statistics are reported on a monthly basis.

	Size	Quality	Value	Momentum	High Vol.	Low Vol.	Liquidity	Market
Mean (bps)	17	3	8	21	33	6	27	5
St.Dev. (bps)	135	145	48	431	428	45	255	167
Sharpe Ratio	0.118	0.020	0.166	0.048	0.078	0.120	0.106	0.028
Skewed T-VaR ₉₉	-4.6%	-4.6%	-1.7%	-13.5%	-13.4%	-1.5%	-8.8%	-6.4%
Skewed T-TVaR ₉₉	-9.3%	-9.5%	-3.2%	-27.8%	-26.9%	-3.1%	-18.1%	-13.2%
Max. Drawdown	23.3%	24.2%	7.6%	54.8%	46.1%	8.2%	34.6%	27.3%
Omega	5.01	3.90	4.98	3.61	3.44	6.35	5.06	3.84
CAPM alpha (bps)	12.5	-0.9	6.9	9.5	22.2	4.3	20.4	
5-Factor alpha (bps)	10.4	-0.8	5.1	7.7	18.6	2.8	19.5	
Tracking Error (bps)	66	55	131	284	276	128	107	
Information Ratio	0.172	-0.032	0.026	0.056	0.104	0.006	0.211	

Table 3.1: Risk and Return measures (excess returns, gross of transaction costs) for all factor portfolios over the period October 2003 - July 2014.

All factor portfolios, with the exception of Quality, have a mean excess return that is at least as high as the excess return of the market portfolio, whose 5bps excess return reflects the market-wide credit risk premium beyond the Term premium. The Low Volatility and Value factor have a mean return similar to the market (8bps and 6bps respectively), whereas the High Volatility and Liquidity factor have mean returns far higher at 33bps and 27bps. Volatility estimates for the factor portfolios can be used to classify Value and Low Volatility as relatively safe factors with lower than market volatility, whereas Liquidity, Momentum and High Volatility can be classed as risky factors, seeing substantially higher volatility estimates. With respect to Sharpe ratios however, factors outperform the market, indicating that the increased risk associated with several factors is more than compensated for in expected excess return.

Figure 3.2 showed that the skewed-t distribution appears to fit the return distributions better and will lead to more conservative VaR measures; this is even more pronounced in the TVaR estimates. The smallest VaR₉₉ is for the Low Volatility at only 1.5% (monthly) and is followed in ascending order by Value, Quality, Size, Liquidity, High Volatility and Momentum, with a VaR of 13.5%.

The highest maximum drawdowns are observed for the Momentum and High Volatility factors, with estimates of 55% and 46% respectively. Relative to the market portfolio, with a drawdown of 27%, the drawdown of Value is very low at only 7%. The only Omega estimates lower than the market estimate of 3.8 are for the Momentum and High Volatility factor at 3.6 and 3.4. These results are not unexpected from a cursory inspection of the empirical histograms and estimated densities since both strategies exhibit positive skewness.

In terms of outperformance versus the market, positive values for alpha are observed, both under the standard-CAPM framework as well as the 5-factor model, indicated by statistically significant non-zero intercepts. These significant results indicate that the higher excess returns are not merely a compensation for taking more systematic risk (CAPM model, estimated by β) or compensation for known equity risk premia (5-factor model).

The relative risk of all single factor portfolios is large considering the market wide credit risk premium (5bps per month); this results in large tracking errors and low information ratios. It is important however, in interpreting the estimated information ratios, to bear in mind that these are based on returns in excess of duration-matched Gilts. The chosen weighting scheme, which merely invests in 10% of the market, effectively applies a zero weight to 90% of the bond universe; with some portfolio turnover added in, the volatility versus the market is understandable. The relatively small sample of bonds causes a higher risk of outperformance.

The risks associated with outperformance re-iterate the need for a relatively long holding period. Even though the sample only spans about 11 years, the performance of the factors is volatile relative to the benchmark. The tracking error is indicative of the relative risk of single portfolio factor investing. This largely corresponds to the general (equity) factor investment literature, where the performance of factor portfolios has been examined over several decades. Whereas factor investing provides excess returns that materialise over long periods of time, over shorter periods of time the performance of factors can be highly cyclical and excess returns can be negative for considerable periods of time. In a research report published by MSCI

(Bender et al., 2013), the performance of six factors is examined over a 25 year period (1988 - 2013). They note how each of the well-documented factors included in this study has experienced periods of at least 2-3 years of under-performance, with the Small Cap (Size) factor experiencing 6 years of under-performance during the 1990s. They conclude that factor investing is no free lunch and the continued existence of excess returns for factors might be a result of the relatively short time horizon (less than several years) a typical market participant (in equities) adheres to; the excess returns cannot be arbitrated away. The period under investigation is relatively short (121 months), and one should be careful about drawing any conclusion regarding the performance of factor portfolios given observations from equity factor outperformance volatility. However, even in the short time period under investigation, the relative risk is apparent.

One can argue that, in order to be successful in investing at factor portfolios, an investor would need to:

- Achieve superior timing mechanism
- Set a sufficiently long time horizon, with an appropriate risk tolerance

Achieving superior timing for the initial investment, with respect to the different factors, might be a viable option for market participants with relatively short investment horizons, but, arguably, will be difficult to realise. One could argue that it is *only* institutional investors that are insensitive to time of entry, specifically due to their long investment horizon and constant investment in the market. The naturally long (intended) investment horizon of insurers and pension funds, typically more than 15 years, appears to be a successful premise to earn factor premia.

Despite having the sufficiently long investment horizon that appears to be ideal to earn factor premia, institutional investors may be partially reliant on external parties, managing (parts of) their portfolio. External fund managers might be reluctant to take advantage of the existence of factor risk premia, despite their clients having a sufficiently long investment horizon. Benchmarked fund managers might be unlikely to pursue a strategy that would see the fund under-perform for prolonged periods of time, especially if compensation is directly tied to outperformance

of the benchmark. Whereas a relatively passive multi-factor approach might reduce the risk of prolonged under-performance, a fund manager arguably has a better chance of outperforming the market on a yearly basis, by adopting a passive strategy to replicate market, seeking out opportunities as they arise (active alpha, or core-satellite approach).

3.4.4 Turnover & Transaction Costs

In a relatively illiquid market, replicating factor portfolios in the corporate bond market might prove impractical if portfolio turnover is such that transaction costs are too high. Trades generally occur for three reasons: maturing of a bond, default-/downgrade of a bond, or the bond fails to meet inclusion criterion for the factor. Section 3.2 explored the extent to which bonds disappear as a result of survivorship bias in the dataset. The methodology for dealing with disappearing bonds by computing a final return value which captures the expected loss is discussed. For the purposes of transaction costs, the expected loss is assumed to include transaction costs. Naturally, only a small proportion of bonds in a portfolio matures in any given month, but the number varies considerably between factor portfolios (see Figure B.2 in Appendix B) where for example, the Quality factor contains bonds with short maturities by definition².

Instead, the purely incremental transaction costs that arise from maintaining factor exposure, i.e. the switching of bonds that fall outside the factor criteria and tolerance level (Figure B.3 in the Appendix B), are highlighted. For this subset of bonds transaction costs are computed, since they arise purely as a result of the factor allocations. Since the dataset contains end-of-day estimates of both Bid- and Ask prices for each bond, Bid-Ask spreads are easily computed. On the day a bond fails to meet the factor inclusion criteria, the quoted Bid-Ask Spread is taken as a direct measure of transaction cost.

Table 3.2 shows how turnover, transaction costs and returns (net and gross) differ between the factor portfolios for the given level of factor tolerance (λ is 40%).

²The replacement of maturing bonds is not included in the transaction cost calculations

	Size	Quality	Value	Momentum	High Vol.	Low Vol.	Liquidity
Turnover (monthly)	0.2%	1.7%	11.1%	42.6%	5.1%	0.7%	6.0%
Return (bps, gross)	17	3	8	21	33	6	27
Return (bps, net)	16	1.3	5.8	-51	26	3.6	21

Table 3.2: Monthly turnover and excess returns (gross and net) over duration-matched Gilts for each of the factor portfolios.

In absolute terms the transaction costs (Bid-Ask Spreads) are considerable for the included Investment Grade bonds during the period 2004-2014. The assumed factor tolerance of 40% seems conservative but still results in excessive turnover for the Momentum factor portfolio, which, as a result, posts substantial negative returns net of transactions costs. All other factors post positive returns net of costs; keeping in mind the 5bps credit premium on the market portfolio, which given the assumptions bears no transactions costs, the Quality and Low Volatility portfolios under-perform the market on a strict excess return basis. Various risk/return characteristics might still favour those factor portfolios however, and the absence of any transaction costs for the market portfolio is an oversimplification.

3.4.4.1 Choice of Factor Tolerance

The results so far are based on the methodology outlined in Section 3.3, using tolerance levels of 40%. This means that a bond will not be switched immediately when it fails to meet the factor criteria (top 10% exposure), but only after it has exceeded the allowed tolerance in addition to the factor definition. Considering the tolerance in the limits, a tolerance of 0% indicates a strategy where bonds in the portfolio at time t , that fall outside the factor criteria at time $t + 1$, are switched immediately; a tolerance of 90% indicates a strategy where bonds in the portfolio at time t will never be switched. Naturally, higher tolerance levels will reduce turnover and lead to lower transaction costs. On the other hand, higher tolerance levels might dilute the factor portfolio leading to decreased exposure to the risk premia, resulting in lower returns. Some optimal tolerance level may exist that maximises the expected return, net of transaction costs. Institutional investors

attempting to earn risk premia may strategically expose (parts) of their portfolio to factors, but maintain a relatively passive approach to investing (buy-and-hold / buy-and-maintain) due to their long time horizon, allowing them to generate income (cash flow) rather than speculate on short term (price) gains. Figure 3.4 shows how changing the tolerance level of the strategy impacts the turnover, transaction costs and returns of each factor portfolio. Note that the Momentum factor has been omitted in Figure 3.4 (right), as net returns are highly negative for most values of λ .

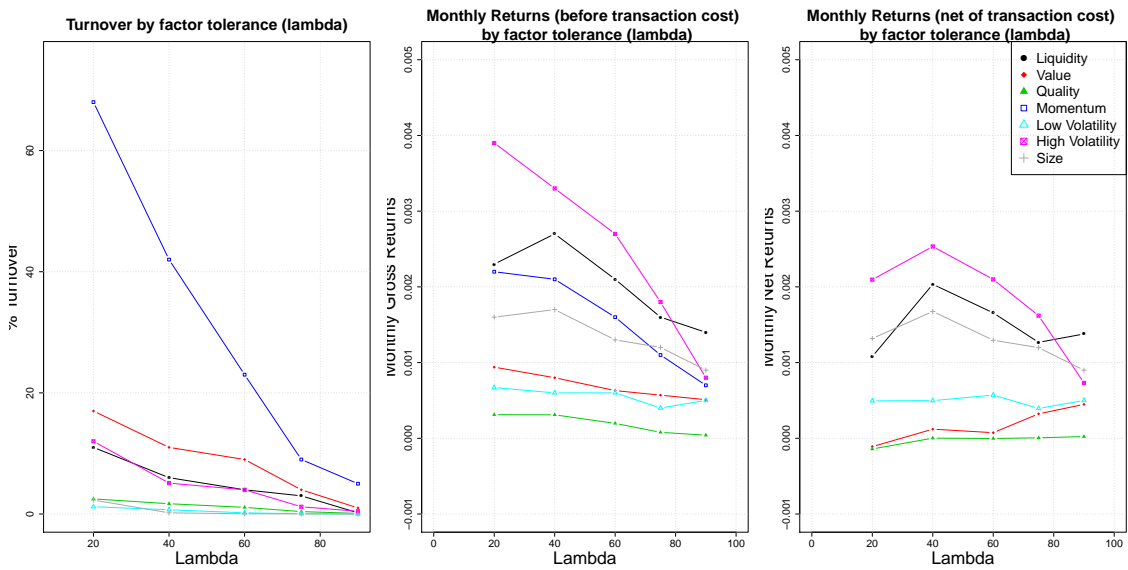


Figure 3.4: For various levels of factor tolerance λ , the turnover (left), gross excess return (middle) and net excess return (right), for each of the seven factor portfolios.

Turnover declines rapidly for high-turnover factors as the tolerance level increases. The difference in monthly turnover for the Momentum, Value, High Volatility and Liquidity factor decreases by approximately 50% when lambda increases from 20% to 60%. Figure 3.4 (middle) shows that factor dilution generally causes excess return to decrease, but not all factors appear to be equally affected. Whereas the difference in gross excess return between low and high values of lambda for the High Volatility and Momentum portfolio is more than 50%, the Liquidity, Value and Size portfolios are far less affected and the Quality and Low Volatility are hardly affected at all; note that the last two factor portfolios saw the lowest turnover. In Figure 3.4 (right) optimal values of the tolerance level may be found; decreased turnover reduces transaction costs whereas factor dilution decreases excess return to some

degree. Firstly, net returns for the Momentum factor are observed to be negative across all tolerance levels. The low yielding factors (Low Volatility and Quality) only saw minor decreases in gross return and very low levels of turnover across tolerance levels and are relatively unaffected by the introduction of transaction costs. The Value factor, also relatively low yielding, saw little change in gross return across tolerance levels, but did see a substantial decrease in turnover, leading to a strictly increasing function of net returns as a function of tolerance levels. The turnover of the Size factor is, understandably, very low, and the net return therefore mimics the gross return series, showing an optimal tolerance level of around 40%. The High Volatility and Liquidity factor have the same turnover function and both appear to arrive at optimal net returns at similar levels as the Size factor.

The pattern of the factor dilution, visible in Figure 3.4 (middle), shows increasing gross returns for the Liquidity factor (to a lesser extent also for the Size factor) when moving from $\lambda = 0.2$ to $\lambda = 0.4$. This appears to suggest there are two ways to take advantage of relatively illiquid bonds: holding to maturity, or, wait for mean reversion of the relative liquidity and sell at a higher price. The data suggest that apparently it may not be optimal, in general, to hold to maturity. Instead mean reversion of the relative liquidity must be sufficiently strong to allow you to crystallise the benefit of the initial high yield at an earlier date when the relative liquidity has closed by a sufficient margin.

3.4.4.2 Annualised Performance Measures

All (risk-adjusted) performance measures and derived statistics are reported using monthly returns rather than their annualised counterparts. While this may make interpretation slightly more difficult (in absolute sense for returns), the difficulty of annualising volatility is avoided for reasons of convenience. Whereas the ‘square-root of time’ is generally accepted as a way to transform volatility from one time unit to another, it is strictly only true under the assumption of independent and identically distributed returns. Whereas the following analysis does not aim to replace the observed monthly statistics with annualised figures, some annualised

figures are reported here to, on the one hand allow for a perhaps straightforward interpretation and comparison, but most importantly to illustrate the process one could go through to annualise returns. Since the returns are dependent over time, the scaling of volatility estimates is less straightforward. However, the dependence structure (Figure B.4 in Appendix B) of the returns can be explicitly modelled and its properties used to appropriately scale volatility. Assume the monthly returns of factor portfolio P follow a first-order autoregressive process with normal innovations (Kaufmann, 2004);

$$P_t = \phi_1 P_{t-1} + \epsilon_t \quad \text{where} \quad \epsilon_t \sim \mathcal{N}(0, \sigma_\epsilon^2) \quad (3.2)$$

where P_t is the monthly return of portfolio P .

In this case, both the 1-month and T-month gains/losses are normally distributed;

$$P_t \sim \mathcal{N}\left(0, \frac{\sigma_\epsilon^2}{(1 - \phi_1)^2}\right) \quad \text{and} \quad \sum_{t=1}^T P_t \sim \mathcal{N}\left(0, \frac{\sigma_\epsilon^2}{(1 - \phi_1)^2} \left(T - 2\phi_1 \frac{1 - \phi_1^T}{1 - \phi_1^2}\right)\right),$$

which gives an expression for the ratio (Q) of the the T-period volatility to the 1-period volatility;

$$Q = \sqrt{\frac{1 + \phi_1}{1 - \phi_1} \left(T - 2\phi_1 \frac{1 - \phi_1^T}{1 - \phi_1^2}\right)}. \quad (3.3)$$

Equation 3.3 illustrates that as $\Phi \rightarrow 0$, $Q \sim \sqrt{T}$. Table 3.3 summarises the annualisation of risk and return statistics after applying the Q-scaling parameter to volatility estimates;

	Size	Quality	Value	Momentum	High Vol.	Low Vol.	Liquidity	Market
Mean month	0.0017	0.0003	0.0008	0.0021	0.0033	0.0006	0.0027	0.0005
Mean year	0.0204	0.0012	0.0096	0.0252	0.0396	0.0072	0.0324	0.0060
StdDev month	0.0135	0.0145	0.0048	0.0431	0.0428	0.0045	0.0255	0.0167
Sharpe month	0.118	0.020	0.166	0.048	0.078	0.120	0.120	0.0278
StdDev year ($\sqrt{12}$)	0.0467	0.0502	0.0166	0.149	0.1482	0.0155	0.0956	0.0579
Sharpe year ($\sqrt{12}$)	0.408	0.692	0.5774	0.1662	0.2701	0.4156	0.3263	0.1037
StdDev year (Q)	0.0827	0.0304	0.0234	0.2255	0.2341	0.0235	0.1495	0.0856
Sharpe year (Q)	0.2465	0.3943	0.4091	0.1117	0.2204	0.2548	0.2086	0.0701

Table 3.3: Annualised volatility estimates using different scaling parameters; scaling under the assumption of independence, and the scaling factor under the assumption of an estimated AR(1) process.

Assuming independent innovations ($Q = \sqrt{12}$) clearly understates the scaling factor and therefore underestimates the volatility on an annual basis for each factor portfolio, including the market portfolio. Even though these methods are a vast improvement of the square root of time rule, they do not come without caution. The scaling parameter in Equation 3.3 assumes a return distribution following an AR(1) and i.i.d errors, but not necessarily a Gaussian error term, ϵ_t , as is assumed in Equation 3.2. This is a rather strong assumption given the empirical data; see Kaufmann (2004) for a more elaborate discussion.

3.4.5 Illiquidity Factor as Liquidity Premium

The previous section illustrates, in great detail, the extent to which investing in factors can lead to more favourable (risk-adjusted) returns, compared to investing in a broad market index. In Table 3.1 we can see that the illiquidity factor returns 27 bps versus 5 bps of the market-wide index when using the standard tolerance factor ($\lambda = 0.4$). When interpreting these figures we need to be aware of the following;

- The factor inclusion is based on one particular proxy, derived in Chapter 2. Earlier analyses indicate that this liquidity proxy is uncorrelated with common bond characteristics and always follows a log-normal distribution. It is important to precisely define the liquidity proxy used in the factor analysis as

the proxy, RBAS, has a different ‘liquidity price’ (γ -coefficient in the models) across rating categories (and days).

- The quoted numbers are all based on returns over duration-matched treasuries, which effectively removes the term premium.
- The liquidity premium, intuitively, is defined differently. Whereas the liquidity premium in Chapter 2 is a theoretical construct, the result of a complex effort to disentangle the credit spread and directly linked to the additional return an investor can earn if the bond is held to maturity, the liquidity premium here represents an active effort to maintain exposure to a particular set of bonds.
- The market-wide portfolio will have some exposure to the illiquidity factor. Arguably, the market-wide index will be of ‘average’ liquidity, where one needs to be aware of the effect a ‘market-cap’ weighted index and an equal weighted index may have as liquidity is often thought to be directly linked to size-measures. In a ‘market-cap’ weighted larger issues naturally carry more weight, implying the index may carry a below (equal-weighted) average liquidity.

The overall estimates for the illiquidity factor returns are an estimate that could be earned, monthly, when investing in illiquid bonds over longer periods of time. Naturally, over the period under study, some months have returned a greater liquidity premium, whereas others would have returned a negative liquidity premium. For instance, the liquidity factor can be built on subsets of the entire bond universe, with Table 3.4 showing basic risk measures for the illiquidity factor, constructed separately for each rating.

	Market Liquidity	AAA	AA	A	BBB
Mean (bps)	27	5	10	22	52
St.Dev. (bps)	255	112	158	228	346
Sharpe Ratio	0.120	0.045	0.063	0.096	0.150

Table 3.4: Risk and return characteristics of Liquidity factors constructed on subsets of the bond universe data.

3.5 Multi-Factor Portfolios

Selecting the right mix of factors should primarily be driven by an institution's objectives and constraints. The reasons for pursuing a multi-factor strategy will differ and could include enhanced risk-adjusted returns or limiting downside risk. With respect to the objectives, different institutions will also have varying beliefs regarding the persistence of factor performance going forward. The constraints institutions face are likely to vary widely; constraints are predominantly determined by the governance structure and risk tolerance/attitude. The stronger the internal governance and the longer the investment horizon, the more likely the institution will have a higher risk tolerance; one could think of constraints related to the size of the managed assets.

In addition to 'asset-side' constraints, some institutional investors will face investment constraints related to the liability side. Unexpected demands for cash might cause forced sales or capital requirements might make certain classes of bonds relatively more or less attractive. These liability-side constraints might be driven by regulation or can be due to more important financial risks such as interest rates risk. Arguably, an insurer is far more interested in pursuing the right interest rate risk strategy across assets and liabilities than an insurer is interested in trying to achieve a slightly higher return on those assets; other considerations are likely to take priority.

Assessing the potential diversification benefits of a multi-factor portfolio, Table 3.5 shows that the excess returns (over duration-matched treasuries) of the factor portfolios and the market portfolio are very strongly correlated ($\rho \geq 0.9$).

Considering the outperformance of the factor portfolios relative to the market portfolio (Table 3.6) however, lower and negative correlations can be seen, suggesting diversification potential. In terms of outperformance there appears to be a division between the safer factors (Value, Size, Quality and Low Volatility) and the higher risk factors (High Volatility, Momentum and Liquidity) that could complement each other well. Intuitively, this is what is expected given the nature of the period under study. The safer factors are far less affected by the financial crisis during 2007-2010,

	Size	Quality	Value	Momentum	High Vol.	Low Vol.	Liquidity	Market
Size	1.00	0.79	0.69	0.85	0.86	0.81	0.91	0.91
Quality		1.00	0.71	0.78	0.81	0.84	0.84	0.94
Value			1.00	0.61	0.63	0.79	0.68	0.81
Momentum				1.00	0.96	0.76	0.96	0.92
High Volatility					1.00	0.72	0.96	0.92
Low Volatility						1.00	0.81	0.90
Liquidity							1.00	0.96
Market								1.00

Table 3.5: Pearson correlation between excess returns over duration-matched treasuries for factor portfolios and the market portfolio.

whereas the higher risk factors would have lost substantial amounts; High Volatility and Momentum saw maximum drawdowns of 46% and 55% respectively, compared to 7% and 8% for the Quality and Low Volatility factors (Table 3.1). Diversification benefits might be overestimated due to the sample period under investigation as a significant portion of the sample period is characterised by severe financial distress which may not be an accurate representation of an ‘average’ period of eleven years.

	Size	Quality	Value	Momentum	High Vol.	Low Vol.	Liquidity
Size	1.00	0.66	0.67	-0.58	-0.55	0.71	-0.41
Quality		1.00	0.94	-0.83	-0.80	0.96	-0.75
Value			1.00	-0.87	-0.86	0.97	-0.82
Momentum				1.00	0.91	-0.84	0.86
High Volatility					1.00	-0.85	0.88
Low Volatility						1.00	-0.79
Liquidity							1.00

Table 3.6: Pearson correlation between the outperformance over the market portfolio for all seven factor portfolios.

A traditional mean-variance optimization (Markowitz, 1991) will lead to undesirable results for heavy-tailed assets (Huisman et al., 1998) and does not account for the spectrum of potential objectives/constraints. For instance, (Huisman et al., 1998) note that mean-variance optimizations lead to an overallocation of capital to assets that exhibit high levels of kurtosis, ignoring the more extreme event risk. Instead, bearing in mind the 5bps premium on the market portfolio, several scenarios are offered. Generally, scenarios are considered that enhance risk-adjusted returns, limit downside risk or limit relative risk. The exact ‘goals’ have been arbitrarily defined (often in reference to the market portfolio), but could be considered a realistic

representation of typical targets for a portfolio (manager) to achieve. See Section 3.4.2 for details on the computation of the risk and return measures used in the scenarios.

- Enhance risk-adjusted returns
 - **A1** Seek maximum excess return keeping Value-at-Risk at market levels
 - **A2** Maximise Sharpe Ratio, subject to a minimum excess return of 25 basis points a month
- Limit downside risk
 - **B1** Seek expected excess return of 15 basis points a month, minimizing Value-at-Risk
 - **B2** Seek expected excess return of 15 basis points a month, minimizing Maximum Drawdown
- Limit relative risk
 - **C** Minimise tracking error keeping volatility at market levels

Each of these strategies is a potential reflection of an institution's objectives in general; the constraints individual institutions might face are far more difficult to replicate. For these strategies all factors are assumed to be persistent in the future and each factor included in the analysis. In addition, only long-only portfolios are considered. The exclusion of short positions in the corporate bond market is realistic due to the costs and constraints (Asquith et al., 2013) associated with borrowing of corporate bonds and prevents superior allocations that are highly impractical or unrealistic.

A Monte Carlo sampling approach is used to arrive at optimal portfolios under each of the scenarios by generating random portfolio weights $w = (w_1, \dots, w_K)$ from a K -dimensional Dirichlet distribution with parameters $(\alpha_1, \dots, \alpha_K)$. This is achieved by drawing K independent samples $Y = (Y_1, \dots, Y_K)$ from Gamma distributions each with density;

$$\text{Gamma}(\alpha_i, 1) = \frac{y_i^{\alpha_i-1} e^{-y_i}}{\Gamma(\alpha_i)},$$

and then setting

$$w_i = \frac{Y_i}{\sum_{j=1}^K Y_j}$$

has w_1, \dots, w_K following a K -dimensional Dirichlet distribution. More specifically, samples are taken from a symmetric Dirichlet distribution, where all parameters $(\alpha_1, \dots, \alpha_K)$ are equal. The single value of α is the concentration parameter for which smaller values lead to a sparser distribution; the larger the values, the more the resulting distribution tends to the uniform distribution. The concentration parameter to 1, corresponding to the uniform distribution on a simplex and condition on weights being less than 60% by omitting some scenarios.

In total 20,000 sets of random multi-factor portfolio weights are drawn and risk and return characteristics of each strategy computed (Figure 3.5). The risk and return characteristics are reported before transaction costs.

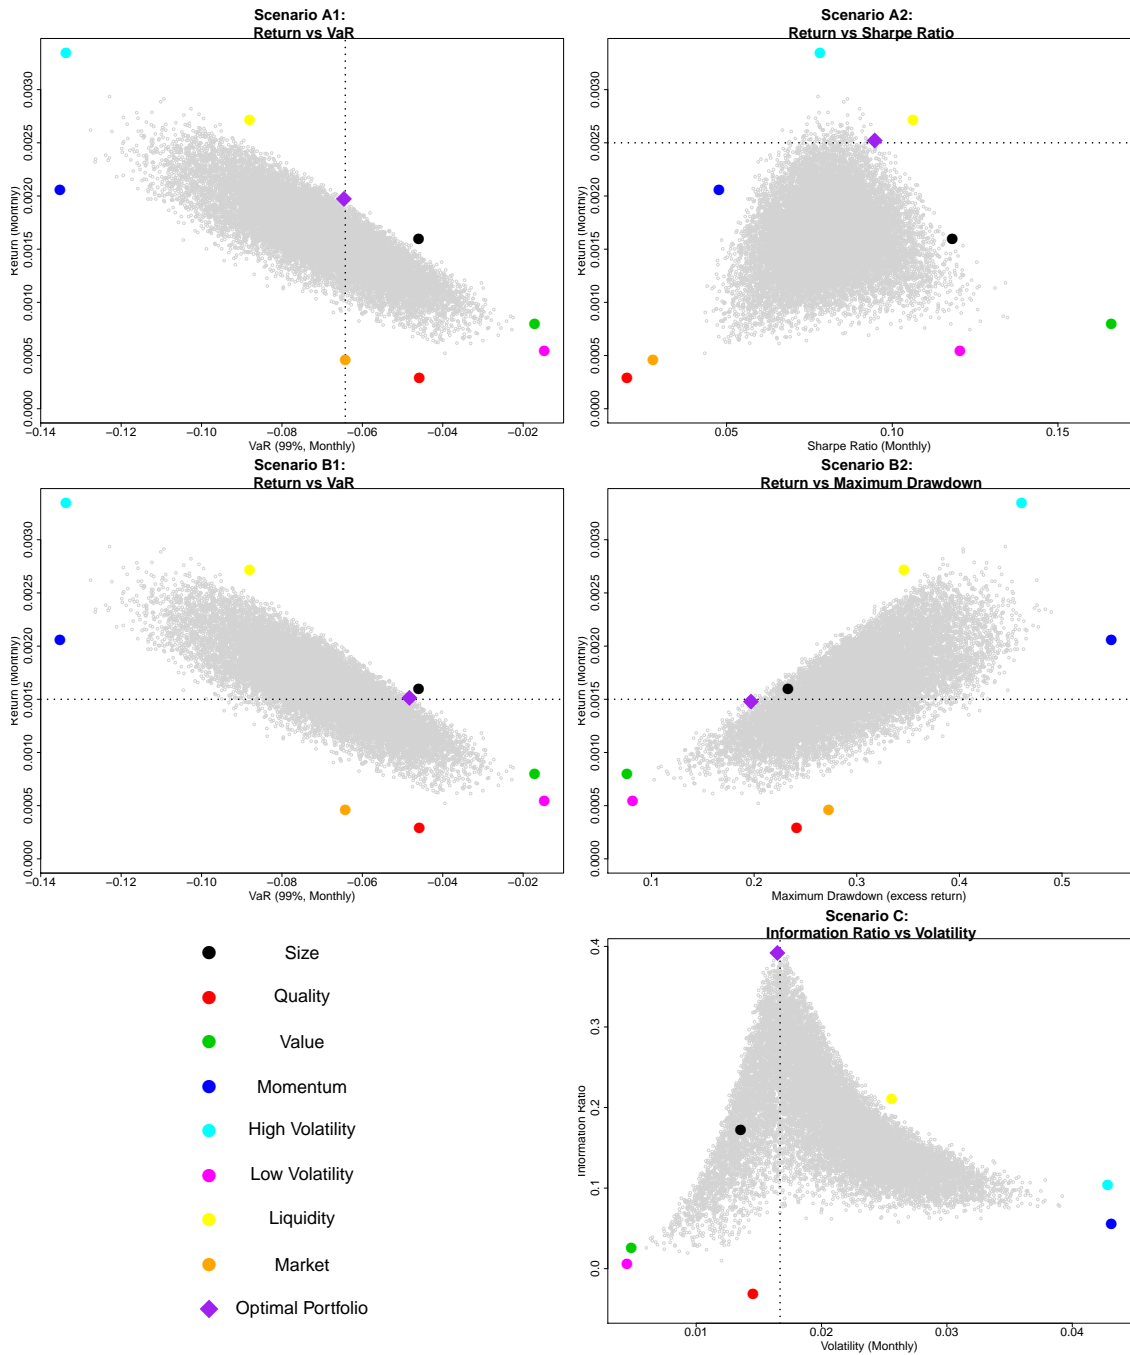


Figure 3.5: Optimisation of multi-factor approach for five scenarios describing an institution's risk/return objective.

The plots in Figure 3.5 show how multi-factor portfolios achieved a return 4 times the return of the market portfolio, for similar levels of VaR (scenario A1), achieved a return 20 basis points above the market portfolio with an increase in Sharpe Ratio of 400% (scenario A2). For a return 10 basis points above the market return, VaR decreased from 8% to 4% in scenario B1 and for the same level of return (10 basis points above the market), maximum drawdown decreased to 16% from 28% for the market (scenario B2). Lastly, keeping the standard deviation similar to the

market portfolio ($\sigma = 0.0167$), Information Ratio (IR) almost double the highest IR of any single-factor portfolio is achieved; the optimised portfolio achieved an IR of 0.395, whereas Size saw the highest IR of any single portfolio at 0.211. Table 3.7 summarises the optimal weights for each strategy as computed using the entire sample period.

	Size	Quality	Value	Momentum	High Volatility	Low Volatility	Liquidity
A1	59.0	1.4	1.7	1.0	5.6	1.0	30.3
A2	20.4	0.2	2.6	3.0	31.0	5.8	37.0
B1	35.4	5.9	10.6	0.9	0.4	22.3	24.5
B2	6.1	4.3	49.8	0.5	6.6	5.7	27.0
C	8.7	0.5	39.2	4.1	2.9	1.0	43.7

Table 3.7: Optimised factor portfolio weights in percentages under each of the strategies.

The optimal portfolio weights as a static measure is only of limited use as it leaves out any information about the variability of these optimal weights or whether near-optimal solutions can be found using different portfolio weights. Before addressing these considerations, two alternative, straightforward, multi-factor weighting schemes are presented. For these two alternative weighting schemes, the relative performance versus the optimal weights (W_O) are noted. Table 3.8 shows a selection of results for an equally weighted portfolio (W_E) and a portfolio with factor weights inversely proportional to its volatility (W_I). Note that not all risk/return metrics are provided for each of the strategies (A1 - C). In Table 3.8, the factor weights of the W_O portfolio varies by strategy while the weighting schemes W_E and W_I are fixed by definition.

Strategy (metric)	W_O	W_E	W_I
A1 (return, bps)	20	16	10
A2 (Sharpe Ratio)	0.093	0.079	0.087
B1 (VaR)	-4.8%	-6.9%	-3.6%
B2 (Drawdown)	19.8%	29.8%	21.4%
C (Information Ratio)	0.39	0.20	0.07

Table 3.8: Comparing two simple weighting schemes, equal weighted (W_E) and inversely proportional to variance (W_I), to the risk and return metric using the optimal portfolio weights (W_O). Each scenario is evaluated here using a difference risk or return metric; the portfolio weights W_O vary for each of the scenarios, whereas the weights for W_E and W_I are fixed by definition.

To see to what extent similar results, that is, similar risk/return characteristics for each of the strategies (A1 - C) can be achieved using portfolio weights different from the optimal weights (W_O), a ‘dissimilarity approach’ is used where a boundary is defined at which results are considered similar (different) and find all sets of portfolio weights that meet this criteria. Then, ultimately, the portfolio weights of all qualifying portfolios are compared with the weights of the optimal portfolio and the least similar portfolio is selected. A portfolio is defined ‘similar’ if the return objective of the portfolio is within 5%³ of the return objective of the optimal portfolio. A straightforward Manhattan distance is subsequently used to measure the distance between the optimal weights (O_1, \dots, O_K) and the alternative weights (A_1, \dots, A_K);

$$d_1(O, A) = \|O - A\|^1 = \sum_{i=1}^K |O_i - A_i|.$$

For all the random portfolios that meet the objective criteria, the portfolio that is least similar, i.e. has the largest Manhattan distance, is selected (Table 3.9).

³A deviation of 5% within the target would, in the example of B1 achieve an excess return over the market of 23.75bps - 26.25 basis point, which a specified target of 25bps.

		Size	Quality	Value	Momentum	High Vol.	Low Vol.	Liquidity
A1	Optimal	59.0%	1.4%	1.7%	1.0%	5.6%	1.0%	30.3%
	Alternative	8.5%	0.7%	2.2%	0.8%	8.6%	30.5%	48.7%
A2	Optimal	20.4%	0.2%	2.6%	3.0%	31.0%	5.8%	37.0%
	Alternative	2.3%	3.5%	8.3%	4.3%	23.8%	4.2%	53.6%
B1	Optimal	35.4%	5.9%	10.6%	0.9%	0.4%	22.3%	24.5%
	Alternative	6.1%	4.3%	49.8%	0.5%	6.6%	5.7%	27.0%
B2	Optimal	6.1%	4.3%	49.8%	0.5%	6.6%	5.7%	27.0%
	Alternative	4.9%	4.7%	49.2%	0.5%	18.8%	10.6%	11.4%
C	Optimal	8.7%	0.5%	39.2%	4.1%	2.9%	1.0%	43.7%
	Alternative	31.6%	0.2%	23.4%	1.7%	10.1%	7.1%	25.9%

Table 3.9: Portfolio weights for the optimal portfolio and the portfolio that meets the similarity criterion, but is least similar, for each of the strategies.

Using rather different portfolio weights, in some strategies, can lead to similar results (within the specified boundaries). This appears to be largely driven by the strong positive correlation (of outperformance) between several factors and the strong negative correlation between others (Table 3.6), effectively offering up alternative factors. On one hand this makes the allocation of factor weights more difficult, as near optimal results can be achieved with portfolio weights that vary, but it is also encouraging as the benefits of diversification and the possibility of meeting a risk/return target are not dependent on an exact specification of portfolio weights, especially given potential uncertainty of beliefs about how these factor portfolios might perform going forward.

Most important, of course, is how a set of portfolio weights will perform out-of-sample, relative to the market and relative to alternative weighting schemes. Given the relatively short period under study, the dataset is not split into a formal training and test set, but the variability of performance is illustrated using a bootstrap-like procedure. Ignoring the serial correlation in the return series (Figure B.4, Appendix B), 10,000 return series are sampled (without replacement), each consisting of 110 months (90% of sample) of data. Using the sampled data, the optimal factor weights

(strategy A1) are determined using the random portfolio approach and the Sharpe ratio (Figure 3.6) is computed for the out-of-sample period with the market portfolio as benchmark. Comparing the out-of-sample performance of the optimal weights, equal weights, weights inversely proportional to individual factor variance, one can gauge the extent to which optimal weights lead to outperformance.




Figure 3.6: Bootstrapped distribution of out-of-sample outperformance versus the market for optimal weights (strategy A1), equal weights and weights inversely proportional to volatility.

Figure 3.6 showed the distribution of the Sharpe Ratio of the out-of-sample return period (versus the market). A Sharpe lower than 0, in this instance, indicates that the alternative weights performed worse than the market portfolio during the out-of-sample period. Since the optimal weights are determined using only 90% of the observations in the sample, this may give *some* indication of future performance. Of the 10,000 simulated return series, the optimal weights outperformed (in terms of Sharpe ratio) the equal weighting scheme in 72%, the inverse weighting in 82% and both in 66%. It remains difficult however to draw conclusions regarding the outperformance of static weights going forward. Arguably, one could allow for a dynamic rebalancing of portfolio weights, in response to changed market perceptions or changing market conditions, but this is beyond the scope of this Chapter,

and arguably very difficult to implement dynamic trading rules that lead to an enhanced performance, out-of-sample. Nevertheless, the bootstrap exercise illustrates how multi-factor portfolios can be employed to decrease the relative risk of factor portfolios and have provided cursory evidence that optimisation of portfolio weights can lead to outperformance, in an attempt to reach pre-defined strategies.

3.5.1 Benefits of Natural Crossing and Reducing Trading Costs

In addition to the potential diversification benefits of constructing multi-factor portfolios, investors should also consider the potential to reduce trading costs through operational efficiency. Combining factor portfolios has the potential to reduce turnover from ‘crossover’ effects. At each month-end, the multi-factor portfolio would be switched to adhere to the separate factors exposure and tolerance targets/rules, and turnover may be reduced as bonds falling out-with one factor’s exposure rules, may be included in another factor portfolio. An example; it could concern a bond that might drop out of the Momentum portfolio due to negative (or low) returns, but included in the Value portfolio as a result. The dynamics of these ‘crossovers’ are not obvious; consider the example of a bond that meets the inclusion criteria for two factor portfolios, on a given day. In the multi-factor portfolio setting this bond is effectively included with an increased weight and the bond falling out-with one of the factor criteria does not lead to a crossover gain. The natural crossovers that are expected, albeit rather infrequently, lead to lower turnover, lower transaction costs and ultimately a higher net return. Section 3.4.4.1 illustrates how turnover is directly related to the tolerance level of the factor exposure criteria; for a tolerance level of 40%, the mean monthly turnovers are 6%, 11%, 2%, 42%, 1%, 5% and 2% for the Liquidity, Value, Quality, Momentum, Low Volatility, High Volatility and Size factor portfolio respectively. If these factor portfolios were replicated separately (and equally weighted), the combined turnover would be 9.5%. Replicated as a multi-factor portfolio, the combined turnover would be lower at 8.7%.

Chapter 4

Model Uncertainty & Parameter Risk in Stochastic Credit Models

4.1 Introduction

The forward projection of prices, returns and yields of credit-risky assets is vital to many financial institutions, whether computing the Value-at-Risk of a portfolio of corporate bonds or performing long-term projections of credit spreads as an insurer. Credit ratings and the corresponding default and transition probabilities over various time horizons, are crucial inputs to stochastic capital models for (corporate) credit portfolios. Another area in which risk models rely heavily on transition probabilities is in credit derivatives, where rating-dependent derivatives are directly impacted by non-default transitions in addition to defaults. Since this risk management function is concerned with rating transitions over short periods of time, typically shorter than the migration matrices published by major credit rating agencies, the correct estimation of these matrices is paramount.

Credit migrations and, in particular, credit defaults, are rare events. The diagonal dominance of commonly published migration matrices is particularly strong for investment grade bonds, making transitions away from the current state rare events, with very few direct defaults over the last few decades. Models that aim to project prices of credit-risky assets into the future frequently rely on the credit ratings as a main risk driver, which, consequently, needs to be projected forward. This

stochastic element to projecting forward ratings relies on known rating dynamics through migration matrices, which are readily available and published by all major credit rating agencies. For the purposes of these models, credit migrations are seen as risk drivers and inputs to a more complex model. The process of credit-risky instruments changing ratings is itself highly complex and far from transparent.

Credit rating agencies assign ratings based on a constant review process that is highly subjective in nature (Goh and Ederington, 1993) and which does not directly relate to any quantitative measure of default. Standard & Poor's (2014) make this explicit in their recent note: 'Credit ratings are not intended as guarantees of credit quality or as exact measures of the probability that a particular issuer or debt issue will default.' Despite observing more downgrades during periods of economic downturn, rating agencies apply a so-called through-the-cycle approach to assigning ratings in an effort to decrease rating volatility; this highlights the fact that ratings are explicitly non-quantitative. Yet, many sophisticated stochastic credit models, as part of commercially exploited Economic Scenario Generators or portfolio simulators, take rating migrations as cardinal or primary risk drivers due to their observable nature.

In addition to applying through-the-cycle principles, credit rating agencies have been criticized for the timeliness of their rating changes (Cheng and Neamtiu, 2009). Rating agencies are cautious in general and attempt to avoid rapid reversals and unnecessary rating volatility (Beaver et al., 2006). The question of timeliness is crucial to the notion of rating drift, the observation that transition probabilities not only depend on the current rating, but also its rating history (see for example Frydman and Schuermann (2008) and Christensen et al. (2004)); after an up-/downgrade, a subsequent change in the same direction is more likely.

In this Chapter the *ignorance risk* when using transition matrices in stochastic credit models for the purposes of risk management or economic capital modelling is investigated. The migration matrix is really the outcome of a modelling process of rating transition events, a model that implicitly makes assumptions about the underlying, real-world process of transitions that are unlikely to hold. These violations

of the Markov properties are well-documented but generally ignored.

Rather than relying on published migration matrices, raw migration data is used in this Chapter, allowing for the construction of alternative models that aim to address some of the known violations of the Markov assumption.

Several approaches to modelling non-Markovian properties of credit ratings are described and competing transition matrices are generated using the same dataset of rating events. A Jarrow, Lando & Turnbull (JLT) credit model (Jarrow et al., 1997) is calibrated, with a stochastic market price of risk (one driver) to spreads of corporate bonds from 2003 - 2014. Simulating from the competing models, with varying underlying migration matrices, accounting for rating momentum, time-inhomogeneity or allowing for statistical uncertainty, it is shown how the distribution of portfolio values is affected and how the distribution of ratings varies at future dates. This approach to evaluating model and parameter risk is pragmatic and practical in nature as ultimately the amount of risk is evaluated using the projections from the baseline model and the projections from the competing models.

Whereas both the phenomena of time-inhomogeneity and rating drift have been well-documented, this is the first study to consider both effects at the same time; for example, to what extent does the business cycle have an effect on the severity of the observed rating drift? More importantly, this study does not intend to demonstrate the extent to which different methodologies or models are exploited to show the existence of non-Markovian properties of rating migrations. Even though a *particular* method of arriving at alternative migration matrices is applied, this is considered ‘as-is’ and the focus lies with the resulting effect on (tail) risk measures of portfolios simulated using a stochastic credit model. In a similar fashion, a relatively straightforward model of stochastic credit spreads is employed where more accurate models may exist; a model with one stochastic driver contains little complexity, but is perfectly suited to the study of risk that arises as a result of alternative credit processes. Not only are the effects of non-Markovian properties of the rating process considered, which are classified as ‘model uncertainty’, but the issue of ‘parameter risk’, interpreted as the risk of not accounting for the statistical

uncertainty in the estimates of the migration matrix, is also investigated. The third ignored risk is perhaps the easiest to address and is the least subject to specification errors or decisions. Credit migration matrices are presented by rating agencies as values, rather than estimates. The estimates are usually the best estimates under the maximum likelihood estimator of the Markov Chain, but carry uncertainty in the same way a quoted correlation coefficient or coefficient from a linear regression carries uncertainty. This uncertainty can be relatively large given the rare nature of the observed events and this uncertainty is quantified and integrated into an alternative model. Lastly, this study briefly considers the sensitivity of parameter choices in the Monte Carlo simulations on the risk measures for all model specifications. This serves two purposes; firstly, it re-iterates the fact that this Chapter attempts to isolate the risk that is ignored in using ready-made migration matrices and does not intend to produce the most accurate Value-at-Risk outcomes, while putting this ignored risk into perspective when some of the other parameters in the model, which are considered to be ‘fixed’, are varied. On the other hand, it illustrates how model uncertainty is generally treated as a brief sensitivity analysis to provide the reader with *some* understanding of the parameter risk.

This Chapter is organized as follows; Section 4.2 reviews some of the underlying principles and discusses the rating process as a Markov process and looks at (industry) standard (stochastic) credit models and extensions, Section 4.3 investigates three ignored risks (rating history, time-inhomogeneity and statistical uncertainty), and Section 4.4 discusses the simulation exercise and the empirical results.

4.2 Review of Modelling Methodologies

This section introduces some of the modelling methodologies that are employed throughout the analysis of migration matrices. In particular the process of moving from observed rating events (upgrades, downgrades and defaults on truncated data) to a real-world and risk-neutral migration matrix representing a Markov Chain, is articulated. The analysis presented here goes one step further and continues to challenge not only the time homogeneity and conditional expectation assumption

that define a Markov process, but also discuss more ‘abstract’ requirements of a Markov process on observed and estimation generator matrices that are engineered from rating events data, or reverse-engineered from published real-world transition matrices. Towards the end of the section, the concepts of real-world and risk-neutral migration matrices, for the purpose of pricing corporate bonds are revisited, with extensive examples using a positive risk premium.

4.2.1 Review of Credit Ratings as a Markov Chain

In this section two general approaches to estimating migration matrices are reviewed; one method is referred to as the ‘discrete’ approach and the other method uses more granular data to construct a Continuous Time Markov Chain (CTMC). The ‘cohort’ approach uses end-of-year rating membership to determine the migration matrix directly. Under the cohort-approach, transition probabilities (q_{ij}) can be readily computed as the total number of rating transitions from rating R_i to R_j ($\sum_{t=1}^T N_{ij}(t, t + 1)$) divided by the number of obligors in R_i at the start of the year (N_i), formally;

$$q_{ij} = \frac{\sum_{t=1}^T N_{ij}(t, t + 1)}{N_i(t)} = \frac{N_{ij}}{N_i}.$$

The likelihood function, $L(P)$, of the Discrete Time Markov Chain (DTMC) comes from K independent multinomial distributions;

$$L(P) = \prod_{i=1}^K \prod_{j=1}^K p_{ij}^{N_{ij}(s)}$$

where $N_{ij}(m)$ is the number transitions from grade i to j during a period of s observations, typically one year.

The main drawback of the cohort approach is that it does not account for rating changes within one time period, but other limitations include the handling of censored data. As a result of those limitations, the cohort approach is known to underestimate rare rating transitions consistently as it does not allow for multiple transition in one time period and is known to estimate rare transitions with high

volatility (Hanson and Schuermann, 2006).

A continuous-time approach to estimate transition probabilities overcomes many of the issues associated with the cohort approach. If λ_{ij} ($i \neq j$) is the instantaneous transition intensity of moving from state R_i to rating R_j and λ_{ij} ($i = j$) is the negative row sum of the off-diagonal elements ($i \neq j$), then Λ_P is the real-world generator matrix of the Continuous-Time Markov Chain $R(t) \in S = 1, 2, \dots, K$.

$$\Lambda_P = \begin{pmatrix} \lambda_{11} & \lambda_{12} & \lambda_{13} & \cdots & \lambda_{1,K} \\ \lambda_{21} & \lambda_{22} & \lambda_{23} & \cdots & \lambda_{2,K} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \lambda_{K-1,1} & \lambda_{K-1,2} & \lambda_{K-1,3} & \cdots & \lambda_{K-1,K} \\ 0 & 0 & 0 & \cdots & 0 \end{pmatrix}$$

where the K^{th} column contains instantaneous probability of default. The absorbing nature of the default state ($\lambda_{KK} = 1$) implies that ratings will move towards the default state in the limit ($\Delta t \rightarrow \infty$). In practical terms this means that a defaulted firm cannot re-emerge as the same entity, but would enter the dataset as a new obligor. In the generator Λ_P , the instantaneous probability of migrating from R_i to R_j between time t and $t + \Delta t$ is $\lambda_{ij}\Delta t$;

$$\lambda_{ij}(\Delta t) \equiv P[R(t + \Delta t) = j | R(t) = i] \geq 0, \forall i, j \in S. \quad (4.1)$$

In order to ensure that a generator matrix produces valid transition matrices with desirable properties for modelling default and credit spreads, several conditions are imposed on the generator (Λ).

1. Off-diagonal elements of Λ should be non-negative:

$$\lambda_{ij} \geq 0 \quad \forall i, j, i \neq j$$

2. All diagonal elements of Λ should be non-positive:

$$\lambda_{ii} \leq 0 \quad \forall i$$

3. The row-sums of Λ should be zero:

$$\sum_{i=1}^K \lambda_{ij} = 0 \quad j = 1, \dots, K$$

4. The default state K is absorbing:

$$\lambda_{Kj} = \begin{cases} 0, & j \neq K \\ 0, & j = K \end{cases}$$

5. Depending on the use of the transition matrix, it might be desirable for the matrix to be stochastically monotonic. This guarantees that the resulting credit spreads are positively monotonic in rating; for example, the credit spread of an AAA bond will always be lower than the spread of an equivalent A or BBB rated bond. This is equivalent to saying that class $i + 1$ is always riskier than class i :

$$\sum_{j \geq k} \lambda_{ij} \leq \sum_{j \geq k} \lambda_{i+1,j} \quad \forall i, k \quad k \neq i + 1.$$

Under the assumption of a known initial state, the likelihood of observations under transitions from i to j at time τ_1 followed by a subsequent transition from j to k at time τ_2 , for each obligor, with rating i ($R_i T$);

$$\begin{aligned} L(\Lambda_P) &= \exp(-\lambda_i(\tau_2 - \tau_1)) \lambda_{ij} \times \exp(-\lambda_j(\tau_3 - \tau_2)) \lambda_{jk} \dots \\ &= \prod_{i=1}^K \prod_{i \neq j} (\lambda_{ij})^{N_{ij}(T)} \exp(-\lambda_i R_i(T)) \end{aligned} \quad (4.2)$$

with log-likelihood;

$$\log L(\Lambda_P) = \sum_{i=1}^K \sum_{i \neq j} \log(q_{ij}) N_{ij}(T) - \sum_{i=1}^K \sum_{i \neq j} q_{ij} R_i(T) \quad (4.3)$$

Therefore the Maximum Likelihood (ML) estimator of Λ_P , with hazard rates $\lambda_{i \neq j}$ is given by;

$$\tilde{\lambda}_{ij} = \frac{N_{ij}(0, T)}{\int_0^T Y_i(s) ds} \quad i \neq j, \quad (4.4)$$

where N_{ij} is the total number of transitions from R_i to R_j , $Y_i(t)$ is the number of obligors rated R_i at time t and $\int_0^T Y_i(s) ds$ is the time spent in rating i over all obligors.

4.2.1.1 Finding a Generator for Public Migration Matrices

The existence of a historical generator matrix is assumed in most academic literature (for example, Figlewski et al. (2012), Hill et al. (2010) or Xing et al. (2012)). The original work by Jarrow et al. (1997) addresses the existence of valid generators, as do Israel et al. (2001) in much more detail. They investigate how and under what circumstances a generator may not exist, and even when one does exist, under what circumstances it may not necessarily be unique. They prove that if the transition matrix is strictly diagonally dominant (i.e. transition probabilities $p_{ii} > 0.5 \quad \forall \quad i$), which historical one-year matrices usually are empirically, and if P is a $N \times N$ transition matrix, then the matrix series \tilde{Q} is defined by;

$$\tilde{Q} = (P - I) - \frac{(P - I)^2}{2} + \frac{(P - I)^3}{3} - \frac{(P - I)^4}{4} + \dots \quad (4.5)$$

is guaranteed to converge and produces an $N \times N$ matrix \tilde{Q} so that $\log(P) = \tilde{Q}$.

Even though this theorem proves to be very powerful due to the empirical observation that historical (one-year) transition matrices, as published by major rating agencies, are always strictly diagonally dominant, it does not guarantee non-negative elements of \tilde{Q} . This becomes problematic since any negative off-diagonal elements in \tilde{Q} will result in negative elements of $P_t = \exp(t\tilde{Q})$ for sufficiently small t , when $t > 0$. Negative transition probabilities in P_t are a clear violation of the Markov property. More specifically, we refer to the classical problem whether a discrete-time Markov process can be embedded in a continuous-time process. This problem is not only related to the modelling of credit risk, and is, for instance, discussed in the context of modelling disability in actuarial health care modelling (Pritchard, 2006). A penalized likelihood solution is used in Bladt and Sørensen (2005) to model the process when a Maximum Likelihood estimator does not exist.

Since the negative off-diagonals in \tilde{Q} tend to be very close to zero, the resulting \tilde{Q} can be made Markovian by simply replacing negative off-diagonals with 0 and adding back the value into the corresponding diagonal to preserve row-sums of zero.

Therefore, using matrix \tilde{Q} , a new generator matrix Q is estimated by setting;

$$q_{ij} = \max(\tilde{q}_{ij}, 0); \quad q_{ii} = \tilde{q}_{ii} + \sum_{j \neq i} \min(\tilde{q}_{ij}, 0) \quad (4.6)$$

Then, row-sums of zero are ensured by setting the diagonal of Q equal to $-\sum_{j \neq i} \tilde{q}_{ij}$;

$$q_{ii} = -\sum_{j \neq i} \tilde{q}_{ij} \quad (4.7)$$

The newly estimated matrix Q will have non-negative off-diagonal entries and maintain zero row-sums. It will however not satisfy $P = \exp(Q)$ exactly. Even if the estimated \tilde{Q} is not a valid generator of P , this does not mean no such generator exists. Even more so, there might be more than one valid generator. For a more extensive discussion of uniqueness of generators, please see Israel et al. (2001).

Rather than investigating the existence and uniqueness of Q generators, the non-existence of valid generators is investigated more easily. The non-existence of a valid generator \tilde{Q} of transition matrix P is conditional on meeting any the following criteria such that;

- (a) $\det(P) \leq 0$ or
- (b) $\det(P) > \prod_i p_{ii}$, or
- (c) *existence of states i and j so that j is accessible from i , but $p_{ij} = 0$,*

then no exact valid generator of P exists.

4.2.1.2 From Published Matrix to Valid Generator

Moody's traditionally publishes data on defaults for various markets internationally. In February 2011 it published its 26th annual default study, in which they update statistics on the default, loss, and rating transitions of corporate bond issuers for 2010 and historically since 1920.

From Moody's Annual Default Study 2011, consider the one-year transition matrix (1980-2011 average), reproduced in Table 4.1.

	Aaa	Aa	A	Baa	Ba	B	Caa	Ca-C	NR	Default
Aaa	87.20	8.21	0.63	0.00	0.03	0.00	0.00	0.00	3.93	0.00
Aa	0.91	84.57	8.43	0.49	0.06	0.02	0.01	0.00	5.48	0.02
A	0.06	2.49	86.07	5.47	0.57	0.11	0.03	0.00	5.13	0.06
Baa	0.04	0.17	4.11	84.87	4.05	0.76	0.16	0.02	5.65	0.17
Ba	0.01	0.05	0.35	5.52	75.75	7.22	0.58	0.07	9.39	1.06
B	0.01	0.03	0.11	0.32	4.58	73.53	5.82	0.59	11.16	3.85
Caa	0.00	0.02	0.02	0.12	0.38	8.70	61.71	3.72	12.00	13.34
Ca-C	0.00	0.00	0.00	0.00	0.40	2.03	9.38	35.46	14.80	37.94

Table 4.1: Migration matrix published by Moody's in their annual default study (Moody's Investor Services, 2011)

The reported table includes rating transitions to the 'Not Rated' category, but no estimates of further default or re-rating are reported (no row for NR). To eliminate this category from the matrix, Jarrow et al. (1997) are followed and the 'Not Rated' category is re-distributed by redefining the transition probabilities from states i to i , when i or $j \neq NR$ using Equation 4.6;

$$p_{ij} = \frac{\text{fraction of firms going from } i \text{ to } j}{\text{fraction of firms going from state } i \text{ to state } \neq NR} \quad (4.8)$$

The resulting transition matrix (denoted by Π) will be used for further analysis. Please note all the elements of the migration matrix below are reported in percentages.

$$\Pi = \begin{pmatrix} 90.77 & 8.54 & 0.66 & 0.00 & 0.03 & 0.00 & 0.00 & 0.00 & 0.00 \\ 0.96 & 89.47 & 8.92 & 0.52 & 0.07 & 0.02 & 0.01 & 0.00 & 0.02 \\ 0.06 & 2.62 & 90.72 & 5.77 & 0.60 & 0.12 & 0.03 & 0.00 & 0.07 \\ 0.04 & 0.18 & 4.36 & 89.95 & 4.30 & 0.80 & 0.17 & 0.02 & 0.18 \\ 0.01 & 0.06 & 0.38 & 6.10 & 83.60 & 7.97 & 0.64 & 0.08 & 1.17 \\ 0.01 & 0.03 & 0.13 & 0.36 & 5.16 & 82.76 & 6.55 & 0.67 & 4.34 \\ 0.00 & 0.02 & 0.02 & 0.13 & 0.44 & 9.88 & 70.12 & 4.23 & 15.16 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.47 & 2.39 & 11.00 & 41.61 & 44.53 \\ 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 0.00 & 100.00 \end{pmatrix}$$

Despite the migration matrix not being strictly diagonally dominant because of element $\Pi_{8,8}$ or $\Pi_{Ca-C,Ca-C} = 0.4161$, the series in Equation 4.5 converges (quickly) and produces $\tilde{\Lambda}$:

$$\tilde{\Lambda} = \begin{pmatrix} -0.0967 & 0.0896 & 0.0069 & 0.0000 & 0.0003 & -0.0001 & 0.0000 & 0.0000 & 0.0000 \\ 0.0101 & -0.1113 & 0.0943 & 0.0055 & 0.0007 & 0.0002 & 0.0001 & 0.0000 & 0.0002 \\ 0.0006 & 0.0275 & -0.0974 & 0.0606 & 0.0063 & 0.0013 & 0.0003 & 0.0000 & 0.0007 \\ 0.0004 & 0.0019 & 0.0459 & -0.1059 & 0.0453 & 0.0084 & 0.0018 & 0.0002 & 0.0019 \\ 0.0001 & 0.0007 & 0.0042 & 0.0666 & -0.1791 & 0.0871 & 0.0070 & 0.0009 & 0.0128 \\ 0.0001 & -0.0002 & 0.0014 & 0.0040 & 0.0566 & -0.1890 & 0.0719 & 0.0074 & 0.0476 \\ 0.0000 & 0.0002 & 0.0002 & 0.0015 & 0.0052 & 0.1174 & -0.3550 & 0.0503 & 0.1801 \\ 0.0000 & 0.0000 & -0.0002 & -0.0001 & 0.0071 & 0.0359 & 0.1652 & -0.8765 & 0.6687 \\ 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 & 0.0000 \end{pmatrix}$$

Using Equations 4.6 and 4.7 to re-distribute the few negative off-diagonal elements in $\tilde{\Lambda}$, the generator is computed. To show that the generator found using the series expansion in Equation 4.5 is more efficient than the original method used by JLT, several distance measures are computed for the NR-adjusted one-year transition matrix published by Moody's and estimates of the migration matrix after transforming to the generator. Estimates using the series expansion follow $\tilde{\Pi} = \exp(\Lambda)$ and the JLT-method follows $\tilde{\Pi}_{JLT} = \exp(\Lambda_{JLT})$, where Λ_{JLT} is computed by the

general formula in their method:

$$\lambda_{ii}^{JLT} = \log(p_{ii}); \quad \lambda_{ij}^{JLT} = \frac{p_{ij} \log(p_{ij})}{p_{ii} - 1} \quad (i \neq j) \quad (4.9)$$

The L1-norm, referred to as the Manhattan distance, and the L2-norm, the Euclidean distance, are used to compare the ‘accuracy’ of the two generator matrices Λ and Λ_{JLT} . The two distance measures are defined as:

1. Manhattan Distance:

$$D_{L1} = \sum_{i=1}^K \sum_{j=1}^K |\Pi_{ij} - Y_{ij}|$$

2. Euclidean Distance:

$$D_{L2} = \sum_{i=1}^K \sum_{j=1}^K (\Pi_{ij} - Y_{ij})^2$$

where Y_{ij} are the best estimates $\tilde{\Pi}$ and $\tilde{\Pi}_{JLT}$ based on generators Λ and Λ_{JLT} respectively. Table 4.2 illustrates that the generator matrix based on the series expansion (Equation 4.6) and a simple redistribution of negative off-diagonals resulted an estimate Π that is many times closer than the generator based on the method outlined in Jarrow et al. (1997).

	Λ	Λ_{JLT}	Factor Gain
L1-norm	0.00257	0.11647	45.30
L2-norm	+0.00000	0.00055	906.31

Table 4.2: Comparison of methods of conversion from published annual migration matrix to generator and back.

4.2.1.3 Historical Data

In the remainder of this Chapter a historical dataset of raw transition events is used to investigate rating transitions and estimations of transition probabilities. The sample of 19,060 corporate bond rating events (S&P) covers the period 1980 to 2002. Following industry conventions to disregard notches brings the number of rating classes to eight, including default. In total, the dataset consists of 19,060 rating events from 10,439 issuers. The total database is geographically divided into North America, United Kingdom, Western Europe & Other and contains data from

debt issued by firms operating in 11 different sectors (Appendix A). Only transition events from North America are included¹, keeping 74% of all rating events. Figure 4.1 illustrates the breakdown of the total dataset (US only) by sector.




Figure 4.1: Sample of rating events (US) broken down over time and sector.

The sampled rating events include a total of 950 defaults, 3340 censored ratings and 3421 ‘NR’ (Not Rated) assignments. Following Altman and Kao (1992), Frydman and Schuermann (2008) and Fei et al. (2012), ‘NR’ assignments are initially treated as just another rating category, but ultimately, transitions to (from) ‘NR’ are assumed not to contain any information about the default risk and are not counted as either downgrades or upgrades. While this may be empirically challenged, following convention, these are transition are disregarded.

A benchmark transition matrix is provided here, based on the Continuous Time Markov Chain approach (equation 4.1), which, for easy comparison and by industry convention, is transformed from the estimated generator matrix into a one-year probability matrix;

$$P(1) = \exp(\Lambda \times 365.25). \quad (4.10)$$

¹Not only does the breakdown of events by domicile change substantially over time, restriction to US events only is convenient in order to link events directly to (US) economic indicators.

From / To	AAA	AA	A	BBB	BB	B	CCC	D
AAA	0.9287	0.0624	0.0073	0.0007	0.0008	0.0001	0.0000	0.0000
AA	0.0061	0.9143	0.0717	0.0061	0.0009	0.0007	0.0001	0.0000
A	0.0008	0.0187	0.9188	0.0549	0.0049	0.0017	0.0001	0.0001
BBB	0.0003	0.0025	0.0404	0.8965	0.0508	0.0076	0.0009	0.0010
BB	0.0003	0.0010	0.0056	0.0494	0.8403	0.0896	0.0076	0.0061
B	0.0000	0.0007	0.0025	0.0049	0.0447	0.8402	0.0584	0.0486
CCC	0.0007	0.0001	0.0029	0.0051	0.0118	0.0797	0.4481	0.4518
D	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	0.0000	1.0000

Table 4.3: One-year CTMC benchmark migration matrix estimated on all rating events from US-based issuers (1980 - 2002).

Because of the multiple transitions a bond can make in a single year, an estimate of the default probability is never zero, despite not observing any transitions directly in the data². The estimated migration matrix in Table 4.3 is the equivalent of the long-term transition matrix published by major rating agencies, typically with a sample period of 30 years (Vazza and Kraemer, 2014).

4.2.2 Review of Default Probability and Credit Spreads

The credit spread (on a zero-coupon bond) can be defined as the difference between the yields to maturity (spot rates) on credit-risky and default free bonds, where i is the rating class at time 0;

$$CS = \frac{1}{T} \log P(0, T) - \frac{1}{T} \log V(0, T, i)$$

The value of a zero coupon bond with maturity T and rating class i , $V(0, T, i)$ can simply be defined as the sum of the discounted future value in case of no default before time T and the discounted future values when the bond does default, recovering only a fraction δ of the face value. The survival probability of rating class i in

²Please note that the default probabilities for, for example, AAA-rated bonds are not *actually* zero. The probabilities are so small that four decimal places is insufficient to display them in Table 4.3.

T years time is described in terms of the generator (Λ);

$$1 - PD_{T,i} = \exp(T\Lambda_{i,K}) \quad (4.11)$$

Next, a straightforward zero coupon bond price can be computed using time-homogeneous default probabilities (real-world), a deterministic (fixed) discount rate r and recovery rate δ_i .

$$V(0, T, i) = \overbrace{V(T, T, i) \exp(-rT)(1 - PD_{T,i})}^{\text{survival claim}} + \overbrace{\delta_i(\exp(-rT)V(T, T, i)(PD_{T,i}))}^{\text{recovery payment}} \quad (4.12)$$

In Figure 4.2 (left), the real-world default probabilities (market price of risk is zero for all i and T) are presented for rating classes $i = 1, \dots, K - 1$ and maturities $0 \leq T \leq 30$. On the right in Figure 4.2, the corresponding Credit Spreads are produced for rating classes AAA (highest) to Ca-C (lowest).




Figure 4.2: Real-world default probabilities (left) and corresponding credit spreads (right) against the term to maturity. The seven curves correspond to seven initial credit ratings; AAA to Ca-C.

In Figure 4.2 (left) a clear link between rating class and default probability is shown; for 'bad' credit ratings (rated B or lower), the structure of the generator illustrates that default will occur quickly or the bond will improve its rating, bringing substantially lower default intensities; these curves are therefore concave. Figure 4.2

(right), with the credit spreads plotted against term to maturity, shows that for high rating classes, the credit spread intercepts the y-axis only at zero. The real-world probability of default is so low for maturities near zero, that it is not priced. This is contrary to what is empirically observed (see for example, Christensen et al., 2004).

So far only transition probabilities and intensities under the real-world probability measure P have been considered. We can introduce a set of transition intensities under the risk-neutral probability measure, Λ_Q , with elements λ_{ij}^Q for $i, j = 1, \dots, k$ where;

$$\begin{aligned} \lambda_{ij}^Q > 0 &\Leftrightarrow \lambda_{ij} \text{ for all } i \neq j \\ \lambda_{ii}^Q &= \sum_{i \neq j} \lambda_{ij}^Q \\ \lambda_{kj}^Q &= 0 \text{ for all } j. \end{aligned}$$

Using both the real-world and risk-neutral transition intensities, we can derive market prices of risk associated with each jump;

$$\mu_{ij} = \begin{cases} \lambda_{ij}^Q / \lambda_{ij} & \text{if } \lambda_{ij} > 0 \\ 1 & \text{otherwise.} \end{cases}$$

From the above it follows that $\mu_{ij} - 1$ can be regarded as the market price of risk for jumps from state i to state j . Similar to Lando (2009), the following is noted with respect to the likely sign and size of μ_{ij} ;

- Bonds that carry default risk are by definition more risky than identical default-free bonds. It is intuitive to assume credit-risky bonds trade at a discount compensating investors for the positive risk premium.
- Diversification within a bond portfolio can significantly reduce the exposure to default uncertainty. In the extreme, perfectly-diversified portfolio, the uncertainty associated with default could be eliminated completely with sufficientl. This equilibrium state implies that default risk is priced under the probability

P (real-world) and no risk premia exist. However, sources of systematic risk factors are present in the market. These risk factors cause dependence between bonds (of various ratings, issuers and other characteristics) and cannot be eliminated by holding a well-diversified portfolio on bonds. The existence of systematic risk factors and risk aversion among investors gives rise to positive risk premia.

- Despite the rating-based modelling approach taken in this Chapter, one could think of each individual bond having its own market price of risk (risk premium). At first, it might seem intuitive to consider two bonds with the same rating, issued by the same corporate entity as having identical market prices of risk. However, there could be numerous reasons for the two bonds to command different risk premia; size of the issue or seniority are two examples.

We follow Jarrow et al. (1997) in showing the effect estimates of μ_{ij} have on the subsequent default probabilities and spreads. Risk-neutral default probabilities are obtained by transforming the generator using scaling parameters, μ_{ij} . Expressed as survival probabilities (as in Equation 4.11);

$$1 - PD_{T,i}^Q = \exp(T\mu_{ij}\Lambda_{i,K}) \quad (4.13)$$

Assuming, for illustration purposes, that for each $i = 1, \dots, n - 1$, $\mu_{ij} - 1 = \mu - 1 > 0$, implying μ_{ij} is independent of credit class, has the effect of speeding up time to exit from state i , but leaving the distribution of transition into different states unaffected. Similar to Figure 4.2, but setting $\mu_{ij} - 1 = 0.4$, the resulting credit default probabilities and zero-coupon credit spreads are plotted in Figure 4.3.

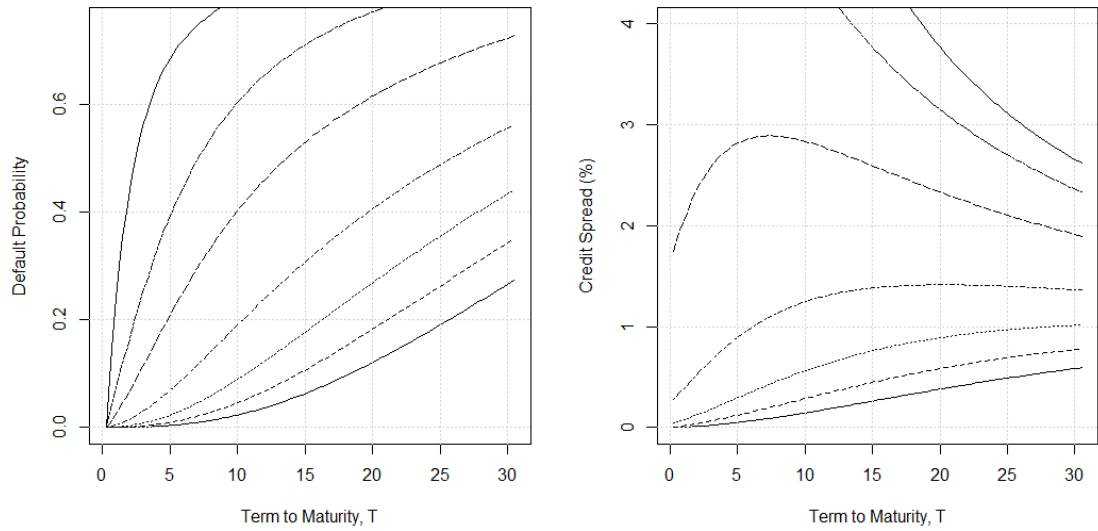


Figure 4.3: Default probabilities (left) and corresponding credit spreads (right) against the term to maturity, introducing a market price of risk $\mu_{ij} - 1 = 0.4$ for all i and t . The seven curves correspond to seven initial credit ratings; AAA to Ca-C.

Comparing Figure 4.2 and Figure 4.3, the effect of introducing a positive market price of risk is clearly observed, where the effect on default probability (left) is most significant for AAA-rated bonds; the same applies to credit spreads where a market price of risk of $\mu_{ij} > 0$ has a large effect. Across rating classes however, the impact of μ_{ij} is observed to be more significant for bonds with shorter maturities. This follows from positive values of μ_{ij} speeding up all times to exit.

Using the resulting credit spread in Figure 4.2 and 4.3, risk premia associated with each rating class and maturity can be computed, given a recovery rate of $\delta_i = 0.4$ and a market price of risk of $\mu_{ij} - 1 = 0.4$. The estimates of the risk premia are easily computed as the difference in credit spread of the credit instrument when using a zero and non-zero market price of risk.




Figure 4.4: Risk Premium as a function of Term to Maturity, where the seven curves correspond to seven initial credit ratings; AAA to Ca-C.

For long-dated bonds, the computed risk premia might look counter-intuitive at first sight (Figure 4.4); bonds perceived to be the riskiest (Ca-C) have the lowest risk premium at long horizons. Over a long period of time (30 years) however, bonds subject to high default probabilities (Ca-C), are almost certain to default; using the generator of real world probabilities (such as the Moody's matrix in Table 4.1), a mere 5% of Ca-C bonds will not have defaulted after 30 year. Therefore, the price of those bonds will be approximately δ times the price of the default-free bond. In that case, a default in the near future has only a small impact on its price. Consequently, the investment can be perceived as far less risky than holding a bond with 30 years to maturity of AAA quality since the impact of re-rating or default on the price of the AAA-rated bond are much greater.

4.3 Ignoring Complexities in the Rating Process

The Markov assumption in credit migrations is often assumed to hold in practice despite the evidence that suggests the rating process is non-Markovian. Ignoring the non-Markovian traits of the rating process is common industry practice and this section describes two examples of such non-Markovian behaviour, modelling these using very simple concepts to illustrate their existence. The other risk under investigation is the risk of ignoring statistical uncertainty of parameter estimates. Rather than the general term ‘model uncertainty’, the term ‘parameter risk’ is used in this instance, as the underlying model itself is not in question nor is the validity of the specification of the model, but the results of the observation that the transition matrix is merely an estimate with a certain degree of uncertainty in its estimates. In summary, this section investigates the effect of ignoring three potential violations of the assumptions that underpin common stochastic credit models that rely on published transition matrices;

- Rating history (non-Markovian process)
- Time variability (non-Markovian process)
- Estimation error / uncertainty.

In the next sections we will see how transition probabilities are dependent on rating history (downgrade history in particular), how transition probabilities are unlikely to be constant throughout time and that published transition matrices carry substantial amounts of estimation uncertainty, which is rarely commented on in rating agency publications.

4.3.1 Rating History

A firm’s rating history does affect future transition probabilities, contrary to the Markov assumption where transition probabilities are conditional on current state only, regardless of its filtration. Rating drift or momentum, the tendency that a rating change is more likely to be followed by a subsequent rating change in the

same direction than suggested under the Markov assumption, is a well-documented phenomenon (Lando and Skødeberg, 2002; Altman and Kao, 1992 and Hamilton and Cantor, 2004). Since rating drift is more pronounced for downgrades (Lando and Skødeberg, 2002; Altman, 1998) than upgrades, rating drift is often narrowed in definition and referred to as downward momentum (Güttler and Raupach, 2010).

Given the existence of downward momentum, the rating process cannot be strictly Markovian. However, following Christensen et al. (2004), returning to a Markovian process by extending the state space with excited states may be a possibility. The transition intensities of the extended state space are then estimated following the usual Continuous Time Markov Chain (CTMC). To capture the tendency for recently downgraded bonds to be downgraded further, a simplified version of the approach taken in Christensen et al. (2004) is used, similar to the approach in Güttler and Raupach (2010). Extending the state space of eight rating classes (including default) to fourteen classes, including excited states;

$$S = AAA, AA, AA^*, A, A^*, BBB, BBB^*, BB, BB^*, B, B^*, CCC, CCC^*, D$$

where * denotes the excited states.

A state is defined as being excited if the rating was previously subjected to a downgrade, where the excited state has infinite memory; it will stay excited until a further rating change. Also note that AAA-rated bonds cannot be excited since the state cannot be reached by a downgrade. At this point it is useful to consider which transitions are feasible under the extended state space; for instance, a transition from A to BBB is not possible. This would constitute a downgrade and subsequently lead to a rating classification of BBB*.

	AAA	AA	AA*	A	A*	BBB	BBB*	BB	BB*	B	B*	CCC	CCC*	D
AAA	X		X		X		X		X		X		X	X
AA	X	X			X		X		X		X		X	X
AA*	X		X		X		X		X		X		X	X
A	X	X		X			X		X		X		X	X
A*	X	X			X		X		X		X		X	X
BBB	X	X		X		X			X		X		X	X
BBB*	X	X		X			X		X		X		X	X
BB	X	X		X		X		X			X		X	X
BB*	X	X		X		X			X		X		X	X
B	X	X		X		X		X		X			X	X
B*	X	X		X		X		X			X		X	X
CCC	X	X		X		X		X		X		X		X
CCC*	X	X		X		X		X		X			X	X
D														X

Table 4.4: Possible migrations under the extended state space are marked with an 'X'.

In the suggested approach, creating extended states where excitement has an infinite memory, it is impossible to move from an excited to a non-excited state within the same rating class (as depicted in Table 4.4). The above model is likely to be the simplest way to model any dependency on rating history. For the purposes of this Chapter, any adjustment for rating drift suffices, but naturally more sophisticated and/or realistic models exist. Perhaps most popular are models that do extend the state space in a similar fashion but allow for a hidden (up-)transition within the rating category with a certain instantaneous intensity (Christensen et al., 2004). Alternatively, the excited states can follow alternative definitions, not necessarily be linked to a historical downgrade, but rather the entire rating history in general.

4.3.1.1 Evidence for Downward Momentum

Before applying the above methodology of extended states to return to a Markovian process, the extent to which the readily-available migration matrices published by major rating agencies hint at non-Markovian properties is investigated, with a particular emphasis on rating momentum.

Published Matrices over time horizons In addition to the (thirty year average) one-year transition matrix, which is the shortest time period for which transition matrices are typically estimated, Moody's also publishes the five-year transition matrix (as a long term average). The published five-year matrix, after correcting for the 'NR'-rated category using Equation 4.8, can be seen in Table 4.5.

	Aaa	Aa	A	Baa	Ba	B	Caa	Ca-C	Default
Aaa	64.09	28.48	6.42	0.43	0.38	0.05	0.05	0.00	0.11
Aa	3.85	61.57	28.00	4.90	0.91	0.28	0.08	0.02	0.40
A	0.25	10.01	65.47	18.67	3.41	1.07	0.22	0.01	0.88
Baa	0.24	1.44	16.45	63.43	11.70	3.73	0.72	0.10	2.19
Ba	0.07	0.27	3.36	19.22	43.54	17.93	2.30	0.18	13.15
B	0.06	0.08	0.48	3.00	11.78	39.67	9.16	1.15	34.63
Caa	0.00	0.00	0.04	1.03	2.99	13.17	16.39	1.86	64.51
Ca-C	0.00	0.00	0.00	0.00	0.00	3.70	3.17	4.53	88.61
Default	0.00	0.00	0.00	0.00	0.00	0.00	0.00	0.00	100.00

Table 4.5: Long term average migration matrix of 5-year transition probabilities published by Moody's in their annual default study (Moody's Investor Services, 2011).

Where the one-year matrix was strictly diagonally dominant (with the exception of the lowest rating category), the five-year matrix is clearly not for several of the lower rating classes. Again, intuitively, one would expect to see lower probabilities across the diagonals as states have more time to jump.

As an estimate of the internal consistency of the one-year and the five-year transition matrix, which will be the basis for evaluating the Markov property, using the generator matrix (Λ_{1Y}) derived from the one-year matrix to estimate the five-year transition matrix by multiplying the one-year matrix by a factor of five using;

$$\tilde{\Pi}_{5Y} = \exp \Lambda_{1Y} \times 5.$$

The matrix D is computed, where D is the element wise difference between

the reported five-year transition matrix (Table 4.5) and the estimated matrix $\tilde{\Pi}_{5Y}$; $D = \Pi_{5Y} - \tilde{\Pi}_{5Y}$. The resulting matrix D shows which elements of the matrix are either overestimated ($D_{ij} > 0$) or underestimated ($D_{ij} < 0$) by the estimator $\tilde{\Pi}_{5t}$.

$$D = \begin{pmatrix} -0.021 & 0.005 & 0.016 & 0.006 & -0.004 & -0.000 & -0.000 & 0.000 & -0.001 \\ -0.008 & -0.016 & 0.020 & 0.011 & 0.001 & -0.003 & -0.001 & -0.000 & -0.004 \\ -0.002 & -0.010 & -0.005 & 0.013 & 0.006 & -0.001 & -0.002 & -0.000 & 0.001 \\ -0.002 & -0.004 & -0.015 & -0.004 & 0.013 & 0.013 & 0.003 & -0.001 & -0.002 \\ -0.001 & -0.003 & -0.004 & -0.012 & 0.015 & 0.021 & 0.017 & -0.002 & -0.052 \\ -0.001 & -0.001 & 0.005 & 0.000 & 0.012 & 0.053 & 0.028 & -0.001 & -0.106 \\ 0.000 & 0.000 & -0.000 & -0.000 & 0.000 & 0.048 & 0.046 & 0.011 & -0.115 \\ 0.000 & 0.000 & 0.000 & 0.000 & 0.010 & 0.023 & 0.038 & -0.025 & -0.046 \\ 0.000 & 0.000 & 0.000 & 0.000 & 0.000 & 0.000 & 0.000 & 0.000 & 0.000 \end{pmatrix}$$

Matrix D illustrates that the discrepancies are relatively small for bonds with a good credit rating (Baa and higher), but non-investment grade bonds show non-negligible differences. The discrepancies are most obvious for the default categories. The fact that the 1-year matrix cannot be converted into the 5-year matrix, is preliminary evidence that the assumptions of a strict Markov process may not be met, which appears to be more pronounced for rating classes of lower credit quality.

Extended States As the original state space is extended to include excited states, the default probabilities of the excited states are hypothesized to be greater than its non-excited counterpart. In an identical bootstrap approach to the benchmark matrix, confidence intervals for the default probabilities are estimated under the extended state space. A comparison of the best estimates against its 99% upper and lower confidence bounds is shown in Table 4.6. A quick inspection of Table 4.6 shows the extent to which default probabilities for the excited and non-excited states differ; there appears to be large differences in default probabilities for low rating classes (B, BB, BBB), no difference for A-rated bonds and an unexpected result for AA-rated bonds, where the default probability for the non-excited state appears to be higher than its excited counterpart. Using the bootstrapped distribution of

default probability estimates, an ANOVA and post-hoc (Tukey) tests are used to formally check the differences. The last column in Table 4.6 and the corresponding level of significance³ confirm what is expected from the cursory inspection of the default rates.

	Lower	Estimate	Upper	Significance
AAA	0.000000	0.000007	0.000013	NA
AA	0.000015	0.000086	0.000138	
AA*	0.000003	0.000048	0.000092	*
A	0.000099	0.000166	0.000224	
A*	0.000084	0.000158	0.000227	
BBB	0.000603	0.001045	0.001397	
BBB*	0.001080	0.002301	0.003330	***
BB	0.005247	0.006409	0.007491	
BB*	0.012532	0.016619	0.020753	***
B	0.036854	0.040490	0.044397	
B*	0.084314	0.095172	0.107225	***
CCC	0.109219	0.146456	0.190627	
CCC*	0.486243	0.521876	0.562614	***

Table 4.6: Best estimates of the default probability for all ratings and states and the corresponding 99% confidence interval.

Table 4.6 shows that the default probabilities of the excited and non-excited states are significantly different for lower-rated classes (CCC, B, BB and BBB), whereas the evidence for high quality classes is less convincing. As an example, Figure 4.5 shows how the bootstrapped estimates of B and B* default probabilities are distributed.

4.3.2 Time Inhomogeneity

The impact of the business cycle or economic indicators on transition probabilities has been studied extensively over the last two decades (see for example Bangia et al. (2002) or Fei et al. (2012)). The underlying premise of this research is that the real economy has a direct impact on a firm's asset value and on its default boundary, as in a Merton model (Merton, 1974). While this makes intuitive sense and is central in the treatment of time-inhomogeneity in this Chapter, it is important to consider that credit rating agencies use a 'Through-The-Cycle' basis, making ratings a somewhat

³* = significance at the 10% level, ** = significance at the 5% level, *** = significance at the 1% level




Figure 4.5: Bootstrapped distribution of 1-year default probabilities for B (left) and B* (right) rated bonds.

relative notion not directly related to actual probability of default (Bar-Isaac and Shapiro, 2013). Closely related to the credit rating agencies' TTC approach, is the potential time lag between movement in economic indicators and rating changes / defaults; Mueller (2008) links activity in the real economy (GDP growth) to credit spreads/rating and finds a lagged response.

A well-known method to illustrate the time-inhomogeneity effects using economic variables is used in this section, but before modelling transition probabilities as a function of economic variables, transitions on an annual basis are explored by fitting the Markov Models on yearly subsets of the data. If the rating process is time homogeneous, annually estimated migration matrices are expected to be the same, subject to noise.

Many approaches have been used to link economic variables to the default process and for the purposes of this Chapter a straightforward approach is used, following Nickell et al. (2000) closely in conditioning the estimation of transition matrices on economic indicators directly. Similar to Nickell et al. (2000), three economic states are specified; 'peak', 'normal' and 'trough', depending on whether the real GDP growth (US) was in the upper, middle two or lower quarter of recorded growth rates

during the sample period (Figure 4.6, right). An alternative model that directly conditions on economic variables is a model similar to Bangia et al. (2002) which conditions on the business cycle directly using the National Bureau of Economic Research (NBER) recession indicator, leaving a recession/non-recession indicator (Figure 4.6, left).

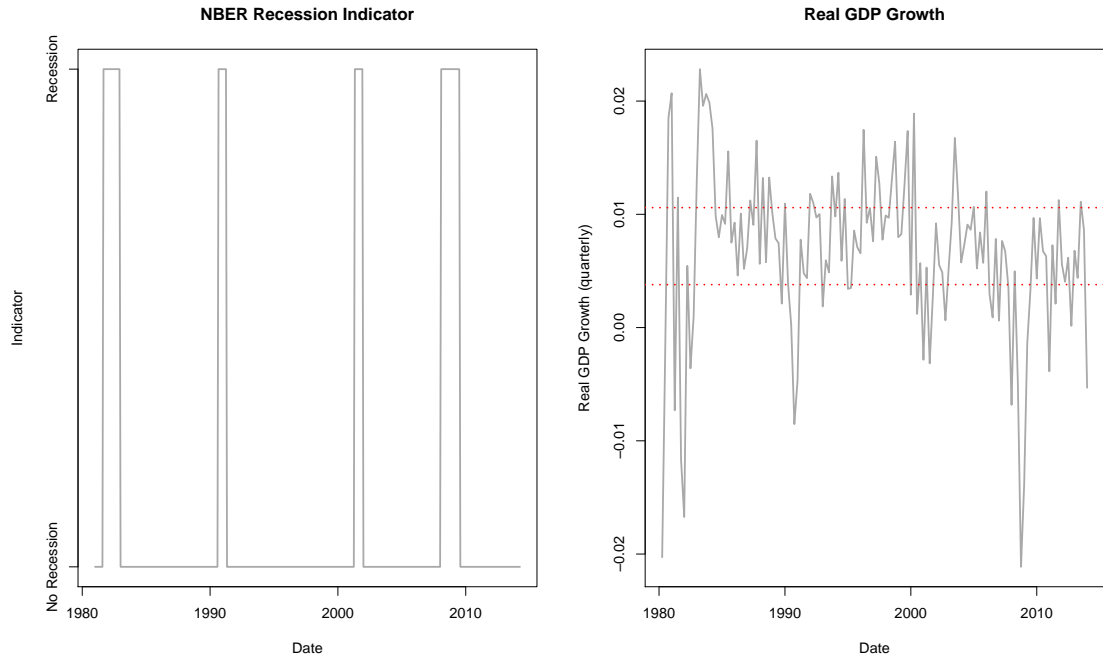


Figure 4.6: Division of sample period by NBER recession indicator (left) and real GDP growth (right).

Conditioning on external economic variables is a convenient way to illustrate how estimated CTMC transitions are different in clearly defined time periods. Ideally however, both the real economy, as well as its relationship with the transition probabilities are treated as a ‘continuous’ process. Estimating a migration matrix using the maximum likelihood approach quickly proves inadequate and unstable when time periods get smaller. To construct migration matrices over very short periods of time, one has to overcome the difficulty of few or no observed rating events in ever smaller periods of time. An attempt to join together discretely estimated matrices over rolling periods of time, so-called ‘dynamic’ transition matrices have been estimated (Berd, 2005). Several alternative approaches exist, which will be briefly highlighted here, but not implemented. The most simplistic, and somewhat ad-hoc solution would be to somehow scale the long-term ‘benchmark’ generator

matrix according to some variable (c_t);

$$\tilde{\Lambda} = c_t \times \Lambda$$

The variable c_t can be a variable believed to be indirectly related to the rating process, or can be directly related to the rating process such as the monthly upgrade/downgrade ratio. Another approach, far less ad-hoc, with very high data and modelling requirements would be to build up the probability of default from first principles on an issuer basis and aggregating over the constituents. The most well-known commercial model is perhaps the Moody's/KMV model, which is capable of producing real-world PD estimates on a monthly basis. This does result, however, in a time series of default probabilities rather than a full transition matrix. A last class of models attempts to build up the transition matrices from scratch using, for instance, hazard models (with frailties), relating transition intensities to observable economic time series (see for example Delloye et al., 2005)

4.3.2.1 Evidence for Time-Inhomogeneity

Before estimating any migration matrices, the number of defaults is explored, as well as the up/downgrade ratio over time. The up/downgrade ratio is defined as $N_{up,t}/N_{down,t}$, where $N_{up,t}$ and $N_{down,t}$ are the number of upgrades and downgrades respectively, in month t . Figure 4.7 shows how these variables vary over time⁴. Generally, this ratio (Figure 4.7, left) is substantially lower than 1, suggesting that bond downgrades are far more common than upgrades.

⁴The number of defaults will be directly related to the size of the relevant S&P bond universe. Without data regarding the size of the universe, this is treated as constant for simplicity.




Figure 4.7: The up/downgrade ratio (left) and the number of defaults (right) vary over the sample period.

Sub-setting the data set by calendar year and subsequently fitting the standard CTMC (an annual ‘benchmark’ matrix), shows how yearly estimates of transition matrices vary. Again, these best estimates are subject to statistical error, similar to all matrices presented in this Chapter. Bootstrapped samples of the annual matrices are not computed, but since the number of observations is split over 22 years, confidence intervals will be much wider.




Figure 4.8: Estimates of default probabilities by rating fitted on annual subsets of the dataset.

Since the state of the economy (GDP growth) is vital in the investigation of time inhomogeneity in this Chapter, these *regimes* are explored further. A time period is classed as a period of ‘low growth’ if the real GDP growth is in the lowest quarter of the growth rates observed during the sample period ($GDP_{\Delta} \leq 0.37\%$) or a ‘peak’ if it is in the upper quarter ($GDP_{\Delta} \geq 1.06\%$). The mean quarterly growth rate is 1.48% for ‘peak’ periods, 0.74% for ‘normal’ periods and -0.31% for low growth periods.

Conditioning on the state of the economy (real GDP growth), patterns clearly emerge. Rather than report the full migration matrix for all three states, Table 4.7 reports the one-year persistence (non-transition probabilities) and the one-year default probability by rating category.

	Low Growth		Normal Growth		High Growth	
	Persistence	Default	Persistence	Default	Persistence	Default
AAA	86.1	0.0	91.7	0.0	96.4	0.0
AA	87.1	0.0	89.5	0.0	95.1	0.0
A	86.4	0.0	88.4	0.0	96.4	0.0
BBB	82.9	0.2	84.9	0.2	95.1	0.0
BB	79.1	1.0	80.4	0.7	91.3	0.3
B	79.0	6.7	80.8	5.9	90.3	2.8
CCC	40.8	50.4	40.0	47.3	55.3	38.1

Table 4.7: Persistence and default probability conditioned on the state of the economy across rating categories.

During ‘peaks’, higher persistence and lower default probabilities across the entire rating spectrum are observed (Table 4.7). Interestingly, greater differences are observed for peaks than for periods of low growth when compared to ‘normal’ periods; this illustrates the importance of modelling time-inhomogeneity as more than just economic downturn. A more restrictive specification of ‘peak’ periods would, most likely, have resulted, similar to a stricter ‘low growth’ definition using the NBER classification, in even more distinct transition probabilities. This would be relatively simple to achieve by defining a peak not as the top quarter of observed GDP growths, but rather the top 10%, for example. It is possible to specify, for instance, 10 buckets of GDP growth, each with an estimated matrix. Figure 4.9 shows how GDP growth appears to be linked to the persistence across ratings and the entire GDP growth spectrum.




Figure 4.9: Rating persistence across all rating categories for ten GDP growth buckets.

If the GDP growth variable were to be extended to ten buckets, a fairly ‘continuous’ variable would be achieved, which, as shown in Figure 4.9, the persistence clearly depends on. Several patterns emerge from Figure 4.9; ratings lower down the quality spectrum generally have lower persistence, persistence generally increases with GDP growth across the entire quality spectrum and sub-setting into ten buckets appears to bring some uncertainty in the estimates. Since persistence takes no direction, low values of persistence could be the result of more mass shifting towards *upgrades*. Far more mass sits in the downgrade portion of the matrices (seen in Table 4.3, as well as the downgrade/upgrade ratio in Figure 4.7), indicating that lower persistence is generally associated with more downgrades and defaults. For the purposes this investigation, allowing for three GDP growth regimes suffices, but this can easily be extended.

4.3.3 Interaction between Rating Drift and the Business Cycle

In addition to investigating rating drift and time-inhomogeneity in separation, the interaction of the two phenomena is studied. Fitting the CTMC with the extended state space, with infinite memory, to periods of recession and periods of

non-recession, as indicated by NBER, provides the following insight. Figure 4.10 displays the difference in default probabilities for excited and non-excited states of the downward-momentum aware matrix fitted previously, as well as those conditioned on NBER recession status.




Figure 4.10: Difference in default probabilities between excited and non-excited counterparts for the CTMC fitted to the entire sample, and the recession/non-recession subsamples. The difference in PD (x-axis) is provided on a log-scale.

Interesting at this point is whether the downward rating drift is stronger during periods of recession and this is investigated by looking at the difference in default probabilities between excited and non-excited states. Figure 4.10 emphasizes once more that downgrade momentum particularly affects lower rated bonds, but also shows that downgrade momentum is more prominent during periods of recession. For CCC rated bonds, the difference in default probability between excited and non-excited state is 27% higher during periods of recessions than non-recession. For B-rated bonds this is 19%, for BB-rated bonds 33% and 23% for BBB-rated bonds. It is important to keep in mind that the combination of the extended state space and the few events that are classed ‘recession’, cause the estimates to have considerable variance; the observed differences are only marginally statistically significant (significant at $\alpha = 5\%$).

4.3.4 Statistical Uncertainty

A crucial part of the investigation into model uncertainty associated with the estimation of migration matrices is statistical uncertainty. The observed (realised) transition events is only one realisation of the true underlying stochastic process. Therefore, an estimated migration matrix, such as the annual matrices published for each calendar year, or the long-term migration matrix (1981-2013) published by Standard & Poor's (Vazza and Kraemer, 2014), are simply an estimates that carry uncertainty. Given the diagonal dominance of migration matrices, one can reasonably expect the largest estimation error to present itself in the far off-diagonals. As default is such a low probability event (1-year matrix) for investment-grade bonds, the observation of rare events (e.g. AAA default) can lead to direct overestimation of the probability; this is even more important when applying a cohort approach. Using a simple parametric bootstrap procedure by taking a random sample of rating histories of all obligors (with replacement), estimates of the $\tilde{N}_{ij}(T)$ and $\tilde{R}_i(T)$ statistics and estimates of generator $\tilde{\Lambda}$ are computed using the ML-estimator in equation (4.4). Lastly, the transition matrix $P(1)$ is derived as $P(1) = \exp(\Lambda \times 365.25)$; 10,000 bootstrap simulations are performed.

For the benchmark migration matrix, ignoring downward momentum and time-inhomogeneity effects, 10,000 bootstrap samples of the generator $\tilde{\Lambda}$ are estimated and, in turn, $P(\tilde{1})$. The 99% confidence intervals for 1-year default probabilities are investigated to gain insight into the statistical error associated with the estimation. Figure 4.11 illustrates the bootstrapped probability distribution of default probabilities for BBB and CCC rated bonds. Table 4.8 summarizes default confidence intervals at the 99% level across ratings.

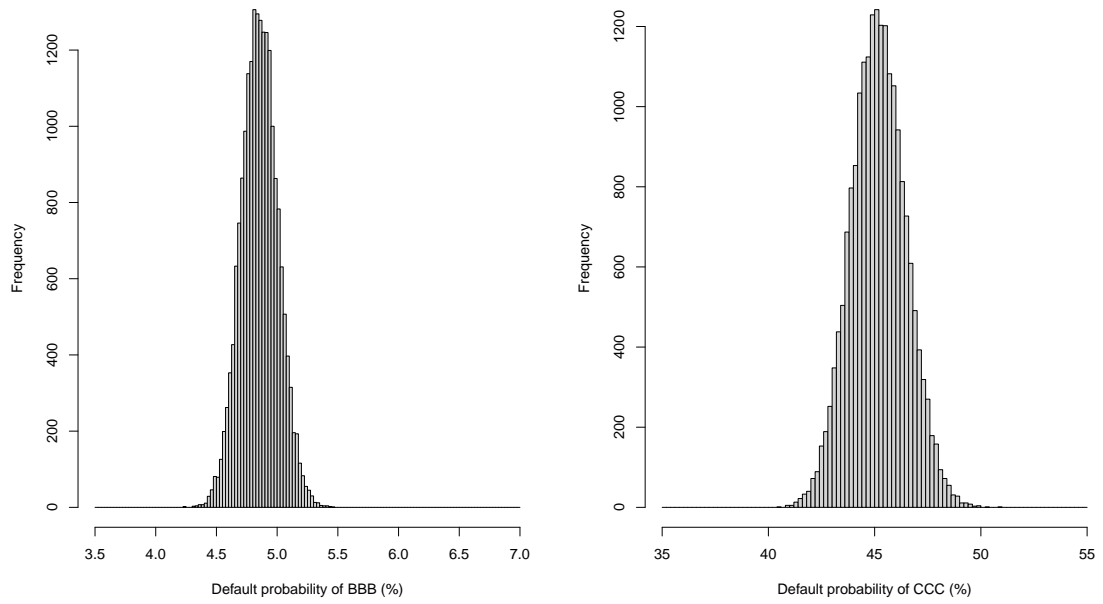


Figure 4.11: Bootstrapped distribution of 1-year default probabilities for BBB (left) and CCC (right) rated bonds.

	Lower	Estimate	Upper
AAA	0.000001	0.000003	0.000006
AA	0.000010	0.000054	0.000086
A	0.000052	0.000084	0.000110
BBB	0.000602	0.001019	0.001386
BB	0.005095	0.006124	0.007120
B	0.045037	0.048554	0.052143
CCC	0.422274	0.451760	0.482689

Table 4.8: Best estimates of the default probability and the corresponding 99% confidence interval.

Table 4.8 shows that the 99% upper bound for the default probability of BBB-rated bonds is 1.36 times the best estimate from Table 4.3; for AA this is 1.59 and for CCC this is 1.06 times. As expected, statistical noise is greater for those matrix elements with the least mass. Even from this initial evidence, the statistical uncertainty appears to be economically significant, even when working with historical long-term estimates based on more than twenty years of data (more than 19,000

rating events of which approximately 1000 are default events).

4.3.5 Stochastic Credit Model & Calibration

Typically, the modelling of credit spread for a specific risky bond features both a jump and a continuous component. The jump part may reflect credit migration and default, i.e. a discontinuous change of credit quality. Meanwhile, credit spreads also exhibit continuous variation so that the spread on a bond of a given credit rating may change even if the risk-free rates remain constant. This may be due to continuous changes in credit quality, stochastic variations in risk premia (for bearing default risk) or liquidity effects.

The JLT credit risk model is based on ratings and thereby centres around the observable change in credit rating rather than changes in a firm's unobservable asset value in a Merton model (Merton, 1974). Using the credit transitions it is fairly straightforward to compute 'break-even' prices (and yields) of credit risky bonds as an expected value of future coupon and principal payments (with pre-determined recovery assumptions). This is calculated using the real-world (cumulative) default probabilities and would be a price at which the market price of risk is assumed to be non-existent and investors are only compensated for the expected credit losses.

As illustrated in Section 4.2.2 when reviewing default probabilities, credit spreads, risk premia and risk-neutral pricing, a Cox-Ingersoll-Ross (CIR) process scales up the generator matrix. This effectively assumes one market price of risk for all rating classes, which follows the same simplification from the theoretical market price of risk attached to each possible transition to a single price of risk, depicted in Figure 4.3. Modelling the market price of risk as a scalar that acts on the generator matrix creates a stochastic credit risk-premium as the difference between risk-neutral (Q) and real-world (P) transition expectations. The CIR process follows the following stochastic differential equation;

$$d\pi_t = \alpha(\mu - \pi_t) + \sigma\sqrt{\pi_t}dZ_t, \quad (4.14)$$

where $\theta \equiv (\alpha, \mu, \sigma)$ are the model parameters. The drift function $\mu(\pi_t, \theta) =$

$\alpha(\mu - \pi_t)$ is linear and mean reverting, resulting in the π -process moving in the direction of μ at speed α . The diffusion function $\sigma^2(\pi_t, \theta) = \pi_t \sigma^2$ is proportional to the risk premium. The risk premium process from Equation 4.14 relates the previously estimated real-world generator to the risk-neutral generator;

$$\Lambda_{Q,t} = \pi_t \Lambda_P.$$

Regarding this stochastic credit process, the following is observed;

- The model will produce strictly positive risk premia, which is the main benefit to modelling the stochastic scalar using a CIR process rather than, for example, a more simplistic Vasicek process. The square root in the variance factor $\sigma\sqrt{\pi_t}$ guarantees strictly positive market prices of risk. Generally, the variance decreases as the process approaches zero reducing the effect of random shocks as the process is dominated by its drift term pushing the process back towards equilibrium.
- The continuous shock component, which comes from the π -process, is perfectly correlated across rating categories. From cursory evidence when exploring the Markit dataset in greater detail in Chapter 2, market-wide credit spreads are observed to be highly correlated across rating categories. In addition to perfect correlation across rating categories, it is worth mentioning that in the simulation exercise, the same risk premium (and shocks to the risk premium) is applied to all individual bonds; the shocks to the stochastic process are modelled through a one-factor model dependency.

Risk neutral probabilities are seen as scaled real-world probabilities in the model which now include compensation for unexpected losses making up the entire credit spread.

4.3.6 Calibrating the Market Price of Risk

Recall that the scaling process π_t from Equation 4.14 follows a CIR process to transform Λ_P to Λ_Q . The π -process is calibrated in two steps. For the period 2003-2015,

information about bond prices, coupon payments/dates, time-to-maturity and credit rating (history) for approximately 1100 bonds are collected on a monthly basis from the Markit iBoxx GBP Investment Grade index constituents. For each month in the dataset the pricing equation is inverted as a function of the risk-neutral default probability and solved for the scalar value π_t that minimises the sum of squared errors between observed prices and prices obtained using the pricing equation, since $\exp(\Lambda_{Q,t}) = \exp(\pi_t \Lambda_P)$. Each bond is given an equal weight in this estimation procedure. Iterating over all months produces a relatively short time series of estimated realisations of $\tilde{\pi}_t$ (139 observations), to which a CIR process can be calibrated using standard methods including least squares regression and maximum likelihood. The calibration procedure is detailed below, where the objective of finding CIR-parameters $\theta \equiv (\alpha, \mu, \sigma)$ that best fit the observed realisations ($\tilde{\pi}_t$) found in the empirical data. In order to quantify the effect of the three risks under investigation, risk neutral credit spreads and bond prices are projected into the future. To investigate the time variability assumption, real world rating dynamics are used as calibrating the CIR-dynamics of the market price of risk would be distractingly difficult due to the existence of distinctly different time periods. Both the downgrade momentum investigation and the statistical uncertainty investigation use different real world dynamics for the rating process (extended state space) and require a different calibration of the parameter set ($\theta \equiv (\alpha, \mu, \sigma)$) of the CIR process that governs the evolution of the market price of risk.

4.3.6.1 Calibration of a CIR process

If $\theta > 0$, $\in \theta$ and $2\alpha\mu \geq \sigma^2$ holds, the CIR process is well defined and has Gamma-distributed marginal distributions. To estimate the parameter vector $\theta \equiv (\alpha, \mu, \sigma)$, transition densities are required. The CIR process has a tractable, closed-form solution. Following the notation in the original publication Cox et al. (1985) on page 391; given r_t , the density of $r_{t+\Delta t}$ is

$$p(r_{t+\Delta t}|r_t; \theta, \Delta t) = ce^{-u-v} \left(\frac{v}{u}\right)^{\frac{q}{2}} I_q(2\sqrt{uv}) \quad (4.15)$$

where

$$c = \frac{2\alpha}{\sigma^2(1 - e^{-\alpha\Delta t})},$$

$$u = cr_t e^{-\alpha\Delta t},$$

$$v = cr_{t+\Delta t},$$

$$q = \frac{2\alpha\mu}{\sigma^2} - 1$$

and $I_q(2\sqrt{uv})$ is a Bessel function of the first kind and order q . The transitional density from Equation 4.15 was originally derived in Feller (1951). From Feller (1951), it is convenient to work with a transformation, $s_{t+\Delta t} = 2cr_{t+\Delta t}$, for which the transition density follows;

$$g(s_{t+\Delta t}|s_t; \theta, \Delta t) = g(2cr_{t+\Delta t}|2cr_t; \theta, \Delta t) = \frac{1}{2c}p(r_{t+\Delta t}|r_t; \theta, \Delta t)$$

which is the non-central χ^2 -distribution with $2q + 2$ degrees of freedom and non-centrality parameter $2u$.

4.3.6.2 Likelihood Function

Parameters estimates of the vector $\theta \equiv (\alpha, \mu, \sigma)$ are obtained from a time series of N observations, with equally spaced time steps of Δt . Using the log-likelihood function with respect to θ

$$\ln L(\theta) = \sum_{i=1}^{N-1} \ln p(r_{t_{i+1}}|r_{t_i}; \theta, \Delta t)$$

from which the log-likelihood function of the CIR process is derived

$$\ln L(\theta) = (N-1) \ln(c) + \sum_{i=1}^{N-1} -u_{t_i} - v_{t_i} + 0.5q \ln\left(\frac{v_{t_{i+1}}}{u_{t_i}}\right) + \ln\{I_q(2\sqrt{u_{t_i}v_{t_{i+1}}})\} \quad (4.16)$$

where

$$u_{t_i} = cr_{t_i}e^{-\alpha\Delta t},$$

$$v_{t_{i+1}} = cr_{t_{i+1}}$$

The maximum likelihood estimates $\hat{\theta}$ are obtained by maximising the log-likelihood function (4.16) over the parameters contained in $\hat{\theta}$:

$$\hat{\theta} \equiv (\hat{\alpha}, \hat{\mu}, \hat{\sigma}) = \underset{\theta}{\operatorname{argmax}} \ln L(\theta) \quad (4.17)$$

4.3.6.3 Initial Estimates using OLS

For the estimates of $\hat{\theta}$ to converge and the log-likelihood function to reach a global maximum, sensible initial estimates are crucial; the likelihood function is particularly flat with respect to α . Brigo and Mercurio (2007) suggest using Ordinary Least Squares (OLS) to arrive at initial estimates. For OLS purposes a discretized version of the SDE in Equation 4.14 is used;

$$r_{t+\Delta t} - r_t = \alpha(\mu - r_t)\Delta t + \sigma\sqrt{r_t}\epsilon_t \quad (4.18)$$

where ϵ_t is a white noise process. In order to perform OLS, equation (4.18) is transformed:

$$\frac{r_{t+\Delta t} - r_t}{\sqrt{r_t}} = \frac{\alpha\mu\Delta t}{\sqrt{r_t}} - \alpha\sqrt{r_t}\Delta t + \sigma\epsilon_t$$

To solve for the initial estimates of the drift parameters $(\hat{\alpha}, \hat{\mu})$ the OLS objective function is minimised;

$$(\hat{\alpha}, \hat{\mu}) = \underset{\alpha, \mu}{\operatorname{argmin}} \sum_{i=1}^{N-1} \left(\frac{r_{t_{i+1}} - r_{t_i}}{\sqrt{r_{t_i}}} - \frac{\alpha\mu\Delta t}{\sqrt{r_{t_i}}} + \alpha\sqrt{r_{t_i}}\Delta t \right)^2$$

The estimate of the diffusion parameter $\hat{\sigma}$ is obtained as the standard deviation of the residuals. Parameters $(\hat{\alpha}, \hat{\mu}, \hat{\sigma})$ are used as initial estimates in Equation (4.17). The calibrated CIR processes π_t (separate process for statistical uncertainty

and downgrade momentum investigation) with parameters $\theta \equiv (\alpha, \mu, \sigma)$ are used to project forward the risk premium; a ‘scalar’ that transforms the real-world generator matrix into the risk-neutral generator.

4.4 Empirical Results using Monte Carlo Simulation

The simulation exercise is used to assess the model uncertainty with respect to the use of migration matrices and highlights three ways in which benchmark stochastic credit models could be misrepresented; time-inhomogeneity, rating drift and statistical uncertainty. Assessing the embedded model uncertainty is through a comparison of the outcomes of a typical simulation exercise of a portfolio of corporate bonds, with a stochastic credit model at heart; the model uncertainty is represented by the different outcomes of the models with respect to Value-at-Risk of the portfolios one year ahead or the rating distribution of the underlying portfolios at various time horizons.

4.4.1 Simulation of Rating Process

To evaluate the impact of different estimation models for migration matrices, the Value-at-Risk (VaR) of a bond portfolio at relatively high confidence levels ($\alpha = 99\%$) is compared over a one-year period (but can easily be used for multi-year projections). The simulation exercise follows an approach similar to simulating credit migrations based on Merton’s model of value of the firm (Merton, 1974).

Obtaining an obligor’s asset value using a one-factor model;

$$X_i = \sqrt{\rho}F + \sqrt{(1 - \rho)}Z_i,$$

where the asset values (X_i) take a common dependence on a systematic risk factor (F) and depend on idiosyncratic shocks (Z_i). Since both F and Z_i are standard normal i.i.d. random variables, ρ is the correlation between changes in asset value among all firms. The dependence model of the obligors’ asset values may be simplis-

tic, and can easily be substituted for more complex or realistic models. One could think of modelling the dependence according to multiple factors, which may reflect drivers such as, for instance, country, industry or currency. In addition one could use a non-Gaussian dependence and employ, for instance, a straightforward t-copula to emphasize a heightened dependence in tail events. Merton's model assumes a firm defaults if its asset value (X_i) drops below some critical value D_i ;

$$I_i = \begin{cases} 1 & X_i < \Phi^{-1}(D_i) \\ 0 & X_i > \Phi^{-1}(D_i) \end{cases}$$

where I_i is the default indicator. The above model is extended to include nine states (all rating classes and default) rather than the default / non-default indicator I_i . By specifying a threshold matrix D with elements d_{ij} , where i refers to the origin state of the rating process, the threshold d_{ij} for non-default states (w) is given by;

$$d_{i,w} = \Phi^{-1}\left(\sum_{w=i}^8 p_{i,w}\right),$$

where $\Phi^{-1}(1) \equiv +\text{inf}$ and therefore $d_{i,1} = +\text{inf}$. Each realisation of the asset value X_i and the threshold matrix D jointly determine the rating one year ahead;

$$S_i = j \quad \text{when} \quad d_{ij} \leq X_i \leq d_{i(j+1)},$$

for the state vector $S = (R_{AAA}, \dots, R_D)$.

If the bond enters the absorbing default state, a stochastic recovery rate is applied, as modelled by a beta distribution, to the market value of the bond (Altman et al., 2004). The asset value generating process, the recovery rate process, the credit risk premium and the risk-free rate are all assumed to be independent. If the year-ahead state is a non-defaulting state, bond prices (B_i) are computed as straightforward expected present values of future cash flows with a random recovery rate in a discrete

setting using an equation similar to 4.12.

$$B_{i,R,t} = \sum_t^T \frac{C_i(1 - PDQ_{R,t}) + C_i(\delta \times PDQ_{R,t})}{(1+r)^t} + \frac{(1 - PDQ_{R,T}) + (\delta_i \times PDQ_{R,T})}{(1+r)^T} \quad (4.19)$$

As stated earlier, the formulated stochastic credit model can be said to be *doubly* stochastic since the model contains both a stochastic jump process (credit rating driver) and a stochastic diffusion element (the CIR process reflecting spread fluctuations, i.e. credit spread risk). In addition to these two drivers of risk, it is important to consider that the model does not consider stochastic interest rates, nor are any of the random processes made to be dependent; the rating, market price of risk and recovery generating process are all constructed to be independent.

To arrive at 1-year risk measures, such as VaR_α estimates, ratings for N bonds in a portfolio are simultaneously simulated forward in time, and project forward the market price of risk scalar to obtain the (risk-neutral) present values for each bond. At the portfolio level this is turned into an aggregate profit/loss (PnL). Repeating the simulation exercise 25,000 times, this approximates the entire PnL distribution.

Every choice regarding the simulation specification determines the results of the simulation exercise, with consequences potentially far greater than the specification errors under study. In order to effectively evaluate the impact the input of migration matrices has, simulation parameters and methods are fixed. However, some insight into the sensitivity of parameter choices (asset correlation, initial rating distribution and portfolio size) on the VaR estimates are presented.

4.4.2 Empirical Results of the Simulated Portfolios

Previous sections have shown how bootstrapped estimates of default probabilities capture statistical noise; Table 4.8 shows that the 99% upper bound for the default probability of BBB is 1.36 times the best estimate from the benchmark matrix. Empirical evidence has been presented that shows that downward momentum is (statistically) significant, with probabilities of default of excited states higher than

the non-excited states in the same rating category. Lastly, approaches to capturing time-varying effects of the rating process show how real-world default probabilities are linked to the real economy.

Based on these observations it is straightforward to hypothesize that VaR estimates using bootstrapped matrices will have higher VaR than the non-bootstrapped counterparts, that the presence of downward momentum leads to higher VaR and that attempts to capture time-inhomogeneity lead to higher VaR estimates. Ultimately, the interest lies with *the extent to which* VaR estimates differ from the benchmark.

The Monte Carlo model is parametrized as follows; simulating a portfolio using 100 bonds ($N = 100$) of 100 notional amount, a maturity date in 6 years ($T = 6$), a coupon rate of 5% paid annually and a fixed annualized risk-free rate of 2% for all maturities. An asset correlation of 0.1998 is used, as estimated by Zeng and Zhang (2001). Both Düllmann et al. (2007) and Lopez (2004) arrive at similar estimates for asset correlations for medium to large US corporations. Recovery rates are simulated from a beta distribution with mean of 0.476 and standard deviation of 0.229 (Altman and Kishore, 1996) in the event of a default.

For the benchmark matrix, the bootstrapped estimates of the benchmark matrix and the migration matrix with an extended state space, the CIR processes are projected forward one year from their September 2015 historical value in monthly time steps⁵. For the simulations challenging the time-homogeneity assumptions, no stochastic scalar is used and instead the focus lies with illustrating how the distribution of rating classes varies from the benchmark using the real world transition dynamics.

Based on 25,000 simulations, the 1-year $\text{VaR}_{\alpha=99\%}$ of the benchmark matrix (Table 4.3) is -4.01%, which will be the value against which the VaR values of competing models based of different migration matrices are compared in Figure 4.12.

⁵Please note that two CIR processes have been calibrated; one based on the benchmark matrix (used in the benchmark en bootstrapped benchmark case) and a process calibrated using the matrix with an extended state space.




Figure 4.12: VaR differences from benchmark matrix for simulation with standard parameters.

Figure 4.12 illustrates clearly how ignoring statistical uncertainty of estimated migration process impacts the simulated risk of the credit risky portfolio. With levels of $VaR_{\alpha=99\%}$ 21% higher than the benchmark, this result is economically significant. Since the statistical uncertainty is straightforward to quantify given the underlying ratings events, and given that running simulations using varying realisations of the same underlying process is straightforward, the results are easily obtained which makes constructing a simulation that accounts for statistical uncertainty a realistic option on a real world commercial setting. The difference in VaR levels is as hypothesized for the rating process that attempts to incorporate rating history under a new Markov process; the VaR of the ‘momentum-aware’ transition model is 49% higher than the benchmark. Using a model based on bootstrapped samples of the ‘momentum-aware’ matrix, thereby accounting for statistical uncertainty in the estimation of the transition matrix with downward momentum, yields, as expected, an ever larger difference in Value-at-Risk of 83%.

For the matrices that have been conditioned on the state of the economy (GDP growth), the regime is allowed to switch on a monthly basis, projecting forward 20 years using the transition matrix in Table 4.9, estimated from the event data in

Figure 4.6 (right).

From / To	High Growth	Normal Growth	Low Growth
High Growth	0.285	0.653	0.062
Normal Growth	0.064	0.847	0.092
Low Growth	0.043	0.736	0.221

Table 4.9: Quarterly transition matrix for three economic regimes, defined by GDP growth.

In each of the simulated months, ratings are projected forward according to the rating migration matrix that corresponds to the regime. Rather than look at Value-at-Risk measures, the focus of the investigation is on the real-world distribution of simulated ratings and how it varies between a model where the benchmark matrix is used and a model where the transition matrix depends on the projected state of the economy. Both simulations start with 1000 bonds of credit quality AA and for simulation years 1, 5, 10 and 20, rating membership is observed for each of the 25,000 simulations. Recording the percentage of bonds in each rating category, for each simulation, gives, at each point in time, a distribution of rating membership for each of the ratings. For instance, using the benchmark model, it could turn out that in year five, a mean percentage of 70% of bonds still labelled A-rated, and, for instance, the 5th quantile is 45% and the 95th quantile is 84%. These numbers are only for illustrative purposes⁶, ultimately the interest lies in the *differences* between the values in the benchmark model and the economy-dependent model, rather than the observed values. Figure 4.13 shows how the difference in 5th quantile and 95th quantile of the rating distributions for the economy-dependent simulation compare to those of the benchmark model, expressed in terms of the benchmark model.

⁶Note that the expectation of the resulting distribution could easily be obtained by taking powers of the 1-year migration matrices.




Figure 4.13: Relative distribution of rating membership at various years ahead, for the benchmark model and the economy dependent model

Figure 4.13 shows how the 5th and 95th quantile of the rating distributions (AAA, BBB and B) are different, at various years ahead, for the benchmark model and the economy dependent model. The results are expressed as a percentage of the benchmark model. For instance, simulation 5 years ahead, the 5th quantile for AAA is only 74% of the benchmark model indicating a lower 5th quantile, which indicates there is the potential for fewer AAA-rated bonds using the economic-dependent model. Figure 4.13 shows that relative to the benchmark, in general, the distribution of rating membership is *wider* for the economy-dependent model. There is both the potential for (far) more, or (far) less rating changes in the economy-dependent model; a result of the randomness in the economic growth that is simulated which can be very different from the average case. Table 4.9 illustrates how GDP growth appears to be highly mean reverting with the highest transition probabilities always towards the medium growth. Evidence of this tendency for mean reversion can be seen in Figure 4.13 as well. The difference between the benchmark model and the economy-dependent model appears to decrease with simulation time, indicating that over longer time period it is rather unlikely that one would experience, say, twenty years of low growth⁷.

⁷As twenty years equals eighty time periods in Table 4.9, this extreme result is expected in

4.4.3 Sensitivity Analysis

A selection of input parameters are varied and used to assess the impact on VaR and VaR differences between estimation methods and the benchmark. The results for varying the asset correlation, initial rating distribution (low/high risk) and portfolio size are summarized in Table 4.10.

	Benchmark	Downward Momentum
Low ρ	64%	103%
Standard ρ	100%	149%
High ρ	213%	246%
Small sample (50)	152%	217%
Standard sample (100)	100%	149%
Large sample (500)	93%	132%
Low risk	74%	102%
Standard risk	100%	149%
High risk	147%	213%

Table 4.10: Value-at-Risk relative to the benchmark matrix with standard model parameters (Benchmark).

As expected, increasing asset correlation, decreasing portfolio size and increasing risk appetite all lead to more higher levels of VaR ($\alpha = 99\%$), where the difference in VaR from the standard parameters can be substantial. The VaR differences arising from changes in some of these model parameters, in particular the very high levels of asset correlation, can be of similar magnitude to VaR differences that arise from model specification and statistical uncertainty. This highlights and leads to an important conclusion that an investigation into model uncertainty during the Monte Carlo simulation stage should not only be concerned with model parameters like asset correlation or recovery rates, but also admit that the rating process is a cardinal input.

Firstly the risk appetite of the portfolio is varied. The low-risk portfolio has an ‘average’ of rating AA, and the high-risk portfolio has an initial rating distribution

only 1 in $\frac{1}{0.221^{80}}$ simulations

with an ‘average’ of BBB; this follows the approach in Jacobson et al. (2006) to assessing the impact of investor’s risk appetite on portfolio return distributions, where the results indicate that a riskier portfolio has a more extreme VaR. The rating distribution (risk appetite) appears to have a stronger effect on the model with a matrix accounting for downward momentum. Intuitively, this makes sense as the difference in default probabilities between excited and non-excited states are most pronounced for ratings lower down the quality spectrum.

Secondly, the portfolio size is varied to represent a small (50 bonds) and large (500 bonds) investor. Jacobson et al. (2006) show that portfolio size is crucial to analysing tail risk in bond portfolios, where they find that larger portfolios have lower VaRs, *ceteris paribus*. This research finds evidence for the same effect and the effect is similar across models. The diversification effect is clearly non-linear, visible even in this small simulation experiment.

Lastly, Güttler and Raupach (2008) are followed and the asset correlation varied to be the 5th (low) or 95th (high) percentile of Zeng and Zhang’s (2001) model, 0.0824 and 0.4331, respectively. Whereas these estimates are more extreme than estimates by Lopez (2004), they clearly provide a best/worst case scenario. Asset correlation is a crucial driver of risk with far-reaching consequences across all models.

Chapter 5

Closing Remarks & Directions for Future Research

5.1 Quantifying the Liquidity Premium on Corporate Bonds

In Chapter 2 a new reduced-form modelling approach to estimating liquidity premia on corporate bonds is proposed that has few data constraints compared to other methodologies using CDS data, structural models or reduced-form models relying on external data for credit risk control variables. The time-varying nature of liquidity premia is demonstrated for various rating categories over an 11-year period capturing a benign financial climate, the financial crisis and more recent years. Liquidity premia are observed, as a proportion of total credit spread, to be bigger for bonds of lower credit quality and emphasize the existence of a distribution of liquidity premia on any given day rather than daily point-estimates. The sign, magnitude and evolution of model parameters provides insight into market dynamics, especially how drivers such as Seniority, Duration Collateralization and Credit Spread changed, broke down or remained stable during the credit crunch. The evolution of parameter estimates in more recent years indicates how market dynamics have recovered, not yet recovered or changed as a result of the crisis.

The model developed in Chapter 2 aims to alleviate some of the difficulties ob-

served in previous modelling efforts that are difficult to calibrate (demonstrated using an analysis of the working paper of the Bank of England that implemented a Leland and Toft model), have extensive data requirements, are difficult to use on a regular basis and provide estimates at only an aggregate level. Using readily available data used to mark to market corporate bond portfolios, a modelling methodology is introduced that produces liquidity premia estimates for individual bonds on a daily basis, without the need to undergo extensive calibration.

Despite the economically intuitive and significant changes in model parameters and outcomes over time, the modelling approach in Chapter 2 defines no explicit time component. Time series analysis could be applied to the daily changes in values for RBAS (on the individual bond level) and its coefficient (for both the individual and aggregate levels). This could provide forward looking estimates of aggregate or bond specific liquidity premia with relatively high frequency (daily/weekly/monthly). Other than, the modelling exercise is conducted entirely cross-sectionally, independent of time.

In addition to the merits of the new modelling approach, several downsides deserve consideration; most importantly, the use of quoted prices and unknown trade/quote sizes are worth discussing. The Over-The-Counter (OTC) nature of the corporate bonds market makes the availability of transaction level data limited. Whereas the Trade Reporting and Compliance Engine (TRACE), where most trades are recorded, is available in the United States, a similar database does not exist in the UK. The use of quotes is second best as the user of the data does not know whether anyone actually acted upon those quotes; one can also wonder how to interpret quotes during times of extreme market distress. Even though the data provider goes to considerable lengths to ensure that the collected quotes from dealer desks are aggregated fairly, other pieces of accompanying information are missing. The average trade size has decreased substantially during the period under study, some of which may be the results of decreased inventories of market makers, which raises the following issue; if the quoted (bid-ask) spreads are recorded for for a 'typical' trade size, then Bid-Ask Spreads are likely to give a flawed picture as quotes

cannot be compared like-for-like over time. Similarly, the methodology is unable to capture the effect of trade size on the liquidity premium and can certainly not conclude whether this effect would vary across bonds with varying degrees of notional amount outstanding.

An area of future work, of particular interest in reference to some of the considerations under the ongoing discussions about the Solvency II regulatory regime with matching assets and liabilities, may focus on the extent investors ‘earn’ the liquidity premium as a function of expected holding period. All estimates of liquidity premia, produced in this thesis and other modelling efforts, relate to the additional expected return when holding the bond to maturity. Even buy-and-hold investors that have the intent to hold to maturity will be faced with an expected holding period that is shorter than maturity due to, for instance, a mandate to sell bonds below BBB, a switch of bonds to lower capital charges, immediate cash needs or a desire to pursue a different risk-bearing strategy. Most importantly, imposing a strict buy-and-hold only policy would make for bad risk management practices. Figure 5.1 is an illustration of what a credit spread decomposition would look like with an explicit time component.




Figure 5.1: Schematic, theoretical decomposition of Credit Spreads, with a time-dependent element.

The expected losses component of the spread (fundamental spread in Solvency II), is split between compensation for losses arising from default and losses from re-ratings. The top component in Figure 5.1 is the accrual of a liquidity premium over time; instantaneous buying and selling of a corporate bond would result in a negative return, of exactly the Bid-Ask Spread. Over time, the liquidity premium (of the acquisition date) is earned. Given that bonds are traded prior to their maturity date for a variety of reasons, one could think of an expected value of the liquidity premium which would be a function of the expected holding period of the asset and the stochastic evolution of the estimates.

5.2 Quantitative Factor Investing in the UK Corporate Bond Market

In Chapter 3 seven factors are reviewed, well-documented in the equity market literature, but for which the credit literature is rather sparse. All factors offer both statistically and economically significant returns beyond the traditional credit risk premium, yet cannot alone be explained by an increased market risk or by exposure to equivalent equity factors in a CAPM or Fama-French framework, respectively. All single-factor portfolios, except Quality, show favourable Sharpe Ratios compared to the market and have significant alphas.

Higher management fees for alternative beta funds are commonly justified due to the (allegedly) partially active nature of the invest strategy, in which an increased turnover can also contribute to the higher transaction costs, which is subsequently passed on to the investors. Due to the illiquid and Over-The-Counter nature of corporate bond trading, transaction costs and turnover are of particular concern. Using a ‘semi-continuous’ approach that allows bonds to be traded (switched) on a monthly basis using a tolerance parameter to control turnover, outperformance appears to be persistent for most strategies using typical values for the tolerance parameter.

Relaxing the fixed holding period offers attractive benefits in favour of a tolerance factor, where portfolio turnover can effectively be controlled. The factor tolerance represents a trade-off between the observed dilution of factor exposures and decreased turnover (transaction costs), with tentative, yet promising suggestions that indeed factor dilution can be observed and there might be an optimal holding strategy that optimizes factor returns net of transaction costs. The tentative results suggest that, depending on the chosen factor, the trade-off between costs and higher returns leads to a suggested optimal holding strategy with medium to high levels of tolerance.

In a multi-factor portfolio context three strategies are specified under which a financial institution might be exploring a multi-factor approach; enhanced risk-

adjusted returns, limited downside risk and limited relative risk. Under each of the strategies and several implementations, opportunities for diversification leading to enhanced returns, reduced risk across a range of metrics, decreased dependency on the economic cycle and reduced relative risk, are achieved.

The results have several implications for institutional investor's managing a large portfolio of corporate bonds. Bearing in mind the cyclical nature of factor performance from equity markets over the last forty years, with the example of more than six years of under-performance of the Small Cap factor in the 1990s, and the volatility of outperformance observed in the smaller time period under study in Chapter 3, the importance of a long-term strategy becomes paramount when it comes to capturing systematic risk premia. The naturally long investment horizon of insurance companies and pension funds seem well-suited to harvesting risk premia that materialise over long time periods. The systematic premia may only exist because the holding period of a typical investor may not be long enough to diversify away the risks at hand.

There are further caveats to consider. Firstly, external money managers might be reluctant to implement factor strategies since tracking errors are large for single-factor portfolios. Combining factor portfolios leads to increased information ratios of up to 0.4, based on returns over duration-matched Gilts. Holding a well-diversified multi-factor portfolio under the proposed methodology would lead to holding up to 650 bonds, with varying weights, at any one time; a large number of bonds to manage, especially given the size of the UK investment grade market. It is very possible that upon further investigation this number can be greatly reduced without affecting the risk and return characteristics of the portfolio. The large tracking error observed of all factor portfolios however is likely the result of having few bonds in factor portfolios, compared to its market-wide benchmark. As the factor portfolios effectively set zero weights to ninety percent of the bonds in the market, alternative strategies to capturing the same premia, using a larger part of the market in order to reduce the volatility of outperformance, might also be of interest.

Secondly, large institutional investors may face additional constraints, unable

to trade bonds to maintain the desired factor exposure. These could include, but are certainly not limited to duration matching of assets and liabilities. The focus of institutional investors would lie with an appropriate cash flow matching and interest rate risks across both sides of the balance sheet, as mismatches in this respect are far more costly than the potential to generate slightly higher returns using a particular investment strategy, that may carry some uncertainty over its outperformance. In addition to these 'top-down' constraints dictated by asset and liability mismatches, institutional investors might have to adhere to internal risk exposure limits (with respect to rating, Duration-Times-Spread or sector) or capital charges that may make individual bonds unsuitable from a portfolio perspective.

Nevertheless, factor investing appears to allow market participants with a sufficiently long time horizon to take advantage of risk premia beyond the credit premium. Depending on beliefs about the persistence of factor returns and subject to individually determined constraints, a multi-factor approach is likely to achieve enhanced risk-adjusted returns, limited downside risk and a decrease in relative risk, regardless of the exact portfolio weights.

Special attention is reserved for the illiquidity factor as an alternative and most importantly intuitive method of deriving liquidity premia. Whereas all other modelling efforts use abstract constructs and complicated models, the intuitive observation of how much additional return is earned on investments in illiquid bonds versus the market, proves useful. Using the, admittedly, arbitrary, relative liquidity proxy developed in Chapter 2 to construct an illiquidity factor, the risk and return characteristics of liquidity premia on several subsets of the investment grade market have been estimated. Whereas this approach appears simple due to the absence of 'models' or any mathematical subjectivity for that matter, it is completely dependent on the chosen criteria for the construction of the factor portfolio.

Future research could focus about the robustness and sensitivity of factor definitions, where the factors constructed in Chapter 3 are (deliberately) simple in definition, there might be far better (smarter) ways to capture some of the mentioned effect. This may be the result of more data, such as balance sheet data,

or different modelling techniques. Rather than relying on traditional equity factors (dumb alpha), it would be interesting to see to what extent smarter algorithms can be designed to forecast the returns of traditional factors for the next twelve months and see to what extent novel data sources can lead to more ‘obscure’ factors. Chapter 3 tried to illustrate to what extent factor returns vary over time and this becomes critical to developing any trading strategy and the need for adequate backtesting becomes apparent. When smart-beta products are concerned, a novel area of research would be the joint study of risk factors across the equity and credit space in strategy design.

5.3 Model Uncertainty & Parameter Risk in Stochastic Credit Models

In Chapter 4 several fairly straightforward, competing (stochastic) credit model specifications are reviewed to capture some well-known non-Markovian properties of the rating process, primarily concerned with the risk of ignorance, that is, the effect of ignoring non-Markovian properties in transition matrices for modelling, for instance, capital requirements. In addition to addressing this case of *model uncertainty*, the effect of statistical error in the estimation of migration matrices is investigated, a fact often overlooked. To study these phenomena effectively, existing literature is followed and straightforward models specified to capture downward momentum by extending the state space conditioned on the state of economy ('Low Growth', 'Normal Growth', 'High Growth') to capture variations over time or take bootstrapped estimates of the migration matrix to capture statistical uncertainty.

Extending the state space to account for downgrade momentum through excited and non-excited states with infinite memory is clearly a restrictive assumption that has previously been modelled in more sophisticated ways, but compelling evidence for downgrade momentum exists nevertheless. Differences in default probabilities between excited and non-excited states are both statistically and economically significant; one-year VaR is 49% higher than the VaR on the benchmark matrix. Some evidence suggests that the downward momentum effect is stronger (between 19% - 33%) during periods of recession than during non-recession periods, but due to small sample size for the extended state space the results are only marginally statistically significant.

Investigating the time-homogeneity assumption of a Markov process, annual default numbers and annual up/downgrade ratios are computed; these are clearly not constant, even after considering that these estimates are subject to statistical uncertainty. Fitting migration matrices to annual subsets of the dataset shows substantial variation in estimated default and persistence probabilities. The convention of subsetting the dataset according to measures of the real economy, focusing on GDP

growth, is followed. The results are certainly economically significant; by specifying either three or ten buckets of GDP growth a clear trend in the persistence and default probability of transition matrices can be seen. Allowing for the economic regime to change on a quarterly basis, evidence exists that the distribution of rating membership is substantially different over a simulation time of twenty years. As the GDP regimes are highly mean reverting, the discrepancy between the benchmark model and the model based on regime switching of GDP growth is less pronounced in long-term simulations.

Accounting for the statistical uncertainty in the migration matrix is relatively straightforward when the underlying transition events are available and accounting for this uncertainty in a simulation exercise gives direct insight into the risk of this ‘parameter’, which appears to be economically significant; Value-at-Risk is 21% higher taking into consideration this uncertainty.

Results indicate that the the studied ‘risk of ignorance’ for simulations using (stochastic) credit models leads to substantial differences, which has direct consequences when estimating, for instance, required capital on credit risky portfolios. For very long term projections of credit risky assets, for the purposes of life insurance for instance, ignoring the parameter uncertainty associated with the migration matrix in favour of a long-term average matrix published by major rating agencies may seem appropriate at first sight. However, whereas the expected value in future time periods does not change, the variance of the distribution certainly does. A similar logic could be made for ignoring the (short/medium) term effects of variations in the economic cycle in very long-term projections. The long-term matrix published by rating agencies is of sufficient length to include all stages of the business cycle and includes periods of financial crises and rapid growth. Therefore, even over longer time horizons the resulting distributions are not too dissimilar, but over shorter time horizons substantial differences arise.

It is worth re-iterating that this cannot be considered a comprehensive treatment of non-Markovian properties of the rating process, nor an attempt to capture rating drift or time-inhomogeneity in the best possible way. While arriving at Value-at-Risk

estimated for hypothetical portfolios of bonds, the actual VaR estimates are of lesser importance than their relative difference, as the assumptions about the portfolio as well the simulation exercise are open for debate. For instance, for high quality bonds the largest sources of risk is arguably not credit risk, but interest rate risk, which is assumed to be fixed and flat. Not to mention more advanced topics that have been studied extensively such as PD - LGD correlation, correlated movements in interest rates and the credit process or alternative dependence structures between the movement of asset values. The differences in VaR estimates when using standard, readily available long-term matrices published by rating agencies versus ‘competing’ matrices are the focus. The relative differences are substantial and this, in turn, will be critical to the capital modelling of *actual* credit-risky portfolios.

To gain a full understanding and appreciation of the model uncertainty, or parameter risks that come with the modelling of credit-risky instruments, a comprehensive review of ‘all the moving parts’ is needed. The aim of Chapter 4 was to discuss in some detail the risks attached to the specification of the rating transition process, which is only one ‘moving part’ of the entire modelling exercise. This comprehensive review could indeed be endless, from varying the dependence modelling of the credit jump process to depend on multiple factors (country, sector) and different dependence models (copulas) to stochastic interest rates or different methods of capturing the stochastic diffusion of the risk premium (credit spread volatility). A comprehensive review would however be unlikely to arrive at any definite conclusions about the model risk of this particular model as a whole, but would provide insight into how different parts of the modelling process are subject to uncertainty and how these model risks may interact.

Appendix A

Appendix to Chapter 2




Figure A.1: Decomposition of credit spread (left) for AAA-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).




Figure A.2: Decomposition of credit spread (left) for AA-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).




Figure A.3: Decomposition of credit spread (left) for BBB-rated bonds of average liquidity into a liquidity and non-liquidity component; Liquidity component of credit spread (middle) in basis points and the liquidity component of credit spread as a proportion of total credit spread (right).

Appendix B

Appendix to Chapter 3

	Aaa	Aa	A	Baa	Ba	B	Caa	Ca-C	NR	Default
Aaa	87.198	8.205	0.631	0.000	0.028	0.002	0.002	0.000	3.933	0.000
Aa	0.908	84.566	8.434	0.492	0.064	0.021	0.008	0.001	5.484	0.021
A	0.057	2.487	86.070	5.475	0.568	0.111	0.032	0.004	5.135	0.062
Baa	0.039	0.172	4.110	84.870	4.054	0.755	0.163	0.017	5.647	0.172
Ba	0.008	0.053	0.348	5.524	75.751	7.220	0.576	0.073	9.387	1.059
B	0.009	0.028	0.113	0.321	4.580	73.526	5.816	0.594	11.159	3.853
Caa	0.000	0.017	0.017	0.116	0.384	8.696	61.708	3.723	12.000	13.339
Ca-C	0.000	0.000	0.000	0.000	0.397	2.034	9.377	35.458	14.797	37.937

Table B.1: Long-term credit migration matrix, estimated by Service in their Annual Default Study of 2014.




Figure B.1: Number of bonds in each of the factor portfolios that disappear from the dataset altogether, but have not matured.




Figure B.2: Number of bonds in each of the factor portfolios that mature in any given month.




Figure B.3: Number of bonds in each of the factor portfolios fall outside the factor exposure criteria, using a tolerance level of $\lambda = 40\%$.




Figure B.4: Sample autocorrelations for the monthly log returns of each factor portfolio.

Appendix C

Appendix to Chapter 4

	North America	UK	Western Europe	World	Total
AeroAuto	9.30	0.30	0.70	1.40	11.70
Construction	4.90	0.10	0.50	0.90	6.40
ConsumerService	11.00	0.40	0.50	1.20	13.10
Energy	4.50	0.10	0.20	0.70	5.50
FinanceInsurance	15.30	1.00	4.00	5.50	25.80
Health	4.60	0.20	0.40	0.40	5.60
Leisure	6.50	0.30	0.20	0.60	7.60
Tech	3.40	0.00	0.20	0.40	4.00
Telecommunications	3.40	0.20	0.40	0.90	4.90
Transportation	3.50	0.20	0.30	0.80	4.80
Utility	7.60	0.70	0.60	1.70	10.60
Total	74.00	3.50	8.00	14.50	100.00

Table C.1: Sector and geographic areas of firms included in the dataset (%).

Bibliography

- Aas, Kjersti and Ingrid Hobæk Haff (2006), “The generalized hyperbolic skew students t-distribution.” *Journal of Financial Econometrics*, 4, 275–309.
- Acerbi, Carlo and Giacomo Scandolo (2008), “Liquidity risk theory and coherent measures of risk.” *Quantitative Finance*, 8, 681–692.
- Altman, Edward, Andrea Resti, and Andrea Sironi (2004), “Default recovery rates in credit risk modelling: a review of the literature and empirical evidence.” *Economic Notes*, 33, 183–208.
- Altman, Edward I (1998), “The importance and subtlety of credit rating migration.” *Journal of Banking & Finance*, 22, 1231–1247.
- Altman, Edward I and Duen Li Kao (1992), “The implications of corporate bond ratings drift.” *Financial Analysts Journal*, 64–75.
- Altman, Edward I and Vellore M Kishore (1996), “Almost everything you wanted to know about recoveries on defaulted bonds.” *Financial Analysts Journal*, 57–64.
- Amihud, Yakov (2002), “Illiquidity and stock returns: cross-section and time-series effects.” *Journal of Financial Markets*, 5, 31–56.
- Amihud, Yakov, Allaudeen Hameed, Wenjin Kang, and Huiping Zhang (2015), “The illiquidity premium: International evidence.” *Journal of Financial Economics*, 117, 350–368.
- Amihud, Yakov and Haim Mendelson (1986), “Asset pricing and the bid-ask spread.” *Journal of Financial Economics*, 17, 223–249.

- Amihud, Yakov, Haim Mendelson, and Lasse Heje Pedersen (2006), *Liquidity and asset prices*. Now Publishers Inc.
- Arora, Navneet, Priyank Gandhi, and Francis A Longstaff (2012), “Counterparty credit risk and the credit default swap market.” *Journal of Financial Economics*, 103, 280–293.
- Asness, Cliff, Andrea Frazzini, and Lasse H Pedersen (2013), “Quality minus junk.” *Available at SSRN 2312432*.
- Asquith, Paul, Andrea S Au, Thomas Covert, and Parag A Pathak (2013), “The market for borrowing corporate bonds.” *Journal of Financial Economics*, 107, 155–182.
- Azzalini, Adelchi and Antonella Capitanio (2003), “Distributions generated by perturbation of symmetry with emphasis on a multivariate skew t-distribution.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 65, 367–389.
- Bachelier, Louis (1900), *Théorie de la spéculation*. Gauthier-Villars.
- Bangia, Anil, Francis X Diebold, André Kronimus, Christian Schagen, and Til Schuermann (2002), “Ratings migration and the business cycle, with application to credit portfolio stress testing.” *Journal of Banking & Finance*, 26, 445–474.
- Bank of England (2007), “Bank of England Financial Stability Report, issue no. 21.” Technical report, Bank of England.
- Banz, Rolf W (1981), “The relationship between return and market value of common stocks.” *Journal of Financial Economics*, 9, 3–18.
- Bao, Jack, Jun Pan, and Jiang Wang (2011), “The illiquidity of corporate bonds.” *The Journal of Finance*, 66, 911–946.
- Bar-Isaac, Heski and Joel Shapiro (2013), “Ratings quality over the business cycle.” *Journal of Financial Economics*, 108, 62–78.

- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny (1998), “A model of investor sentiment.” *Journal of Financial Economics*, 49, 307–343.
- Basu, Sanjoy (1977), “Investment performance of common stocks in relation to their price-earnings ratios: A test of the efficient market hypothesis.” *The Journal of Finance*, 32, 663–682.
- Beaver, William H, Catherine Shakespeare, and Mark T Soliman (2006), “Differential properties in the ratings of certified versus non-certified bond-rating agencies.” *Journal of Accounting and Economics*, 42, 303–334.
- Bender, Jennifer, Remy Briand, Dimitris Melas, and Raman Aylur Subramanian (2013), “Foundations of factor investing.” Technical report, MSCI.
- Benos, Alexandros and George Papanastasopoulos (2007), “Extending the merton model: A hybrid approach to assessing credit quality.” *Mathematical and Computer Modelling*, 46, 47–68.
- Benston, George J and Robert L Hagerman (1974), “Determinants of bid-asked spreads in the over-the-counter market.” *Journal of Financial Economics*, 1, 353–364.
- Berd, Arthur M (2005), “Dynamic estimation of credit rating transition probabilities.” Available at SSRN 842664.
- Berkowitz, Jeremy and James OBrien (2002), “How accurate are value-at-risk models at commercial banks?” *The Journal of Finance*, 57, 1093–1111.
- Black, Fischer and John C Cox (1976), “Valuing corporate securities: Some effects of bond indenture provisions.” *The Journal of Finance*, 31, 351–367.
- Black, Fischer and Myron Scholes (1973), “The pricing of options and corporate liabilities.” *The Journal of Political Economy*, 637–654.
- Blackrock (2015), “The evolution of active investing: Finding big alpha in big data.” Technical report, Blackrock.

- Bladt, Mogens and Michael Sørensen (2005), “Statistical inference for discretely observed markov jump processes.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67, 395–410.
- Bohn, Jeffrey R (2000), “A survey of contingent-claims approaches to risky debt valuation.” *The Journal of Risk Finance*, 1, 53–70.
- Bongaerts, Dion, Frank De Jong, and Joost Driessen (2011), “Derivative pricing with liquidity risk: Theory and evidence from the credit default swap market.” *The Journal of Finance*, 66, 203–240.
- Breger, Ludovic and Darren Stovel (2004), “Agency ratings in the pfandbrief market.” *The Journal of Portfolio Management*, 30, 239–243.
- Brennan, Michael, Sahn-Wook Huh, and Avanidhar Subrahmanyam (2013), “An analysis of the amihud illiquidity premium.” *Review of Asset Pricing Studies*, 3, 133–176.
- Brennan, Michael J, Tarun Chordia, Avanidhar Subrahmanyam, and Qing Tong (2012), “Sell-order liquidity and the cross-section of expected stock returns.” *Journal of Financial Economics*, 105, 523–541.
- Brennan, Michael J and Avanidhar Subrahmanyam (1996), “Market microstructure and asset pricing: On the compensation for illiquidity in stock returns.” *Journal of Financial Economics*, 41, 441–464.
- Brigo, Damiano and Fabio Mercurio (2007), *Interest rate models-theory and practice: with smile, inflation and credit*. Springer Science & Business Media.
- Brunnermeier, Markus K and Lasse Heje Pedersen (2009), “Market liquidity and funding liquidity.” *Review of Financial studies*, 22, 2201–2238.
- Cai, Fang, Song Han, and Dan Li (2012), “Institutional herding in the corporate bond market.” *FRB International Finance Discussion Paper*.
- Carhart, Mark M (1997), “On persistence in mutual fund performance.” *The Journal of Finance*, 52, 57–82.

- Cetin, Umut, Robert Jarrow, Philip Protter, and Mitch Warachka (2006), "Pricing options in an extended black scholes economy with illiquidity: Theory and empirical evidence." *Review of Financial Studies*, 19, 493–529.
- Chen, Long, David A Lesmond, and Jason Wei (2007), "Corporate yield spreads and bond liquidity." *The Journal of Finance*, 62, 119–149.
- Chen, Ren-RAW, Xiaolin Cheng, and Liuren Wu (2005), "Dynamic interactions between interest rate, credit, and liquidity risks: Theory and evidence from the term structure of credit default swap spreads." *Available at SSRN: 779445*.
- Cheng, Mei and Monica Neamtiu (2009), "An empirical analysis of changes in credit rating properties: Timeliness, accuracy and volatility." *Journal of Accounting and Economics*, 47, 108–130.
- Christensen, Jens HE, Ernst Hansen, and David Lando (2004), "Confidence sets for continuous-time rating transition probabilities." *Journal of Banking & Finance*, 28, 2575–2602.
- Churm, Rohan and Nikolaos Panigirtzoglou (2005), "Decomposing credit spreads." *Bank of England Working Paper*.
- Clarke, Roger G, Harindra De Silva, and Steven Thorley (2006), "Minimum-variance portfolios in the us equity market." *The Journal of Portfolio Management*, 33, 10–24.
- Cochrane, John H (2005), "Asset pricing program review: Liquidity, trading and asset prices." *NBER Reporter*.
- Collin-Dufresne, Pierre, Robert S Goldstein, and J Spencer Martin (2001), "The determinants of credit spread changes." *The Journal of Finance*, 56, 2177–2207.
- Copeland, Thomas E and Dan Galai (1983), "Information effects on the bid-ask spread." *The Journal of Finance*, 38, 1457–1469.
- Correia, Maria, Scott Richardson, and İrem Tuna (2012), "Value investing in credit markets." *Review of Accounting Studies*, 17, 572–609.

- Cox, John C, Jonathan E Ingersoll Jr, and Stephen A Ross (1985), “A theory of the term structure of interest rates.” *Econometrica: Journal of the Econometric Society*, 385–407.
- Daniel, Kent and Sheridan Titman (1997), “Evidence on the characteristics of cross sectional variation in stock returns.” *The Journal of Finance*, 52, 1–33.
- Danielsson, Jon, Kevin R James, Marcela Valenzuela, and Ilknur Zer (2015), “Model risk of risk models.” *FEDS Working Paper*.
- de Carvalho, Raul Leote, Patrick Dugnonne, Xiao Lu, and Pierre Moulin (2014), “Low-risk anomalies in global fixed income: Evidence from major broad markets.” *The Journal of Fixed Income*, 23, 51–70.
- Dechow, Patricia, Weili Ge, and Catherine Schrand (2010), “Understanding earnings quality: A review of the proxies, their determinants and their consequences.” *Journal of Accounting and Economics*, 50, 344–401.
- Delloye, Martin, Jean-David Fermanian, and Mohammed Sbai (2005), “Estimation of a reduced-form credit portfolio model and extensions to dynamic frailties.” *Preliminary Version, BNP-Paribas*.
- DeMiguel, Victor, Lorenzo Garlappi, and Raman Uppal (2009), “Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy?” *Review of Financial Studies*, 22, 1915–1953.
- Demsetz, Harold (1968), “The cost of transacting.” *The Quarterly Journal of Economics*, 81, 33–53.
- Deuskar, Prachi, Anurag Gupta, and Marti G Subrahmanyam (2011), “Liquidity effect in otc options markets: Premium or discount?” *Journal of Financial Markets*, 14, 127–160.
- Dick-Nielsen, Jens, Peter Feldhütter, and David Lando (2012), “Corporate bond liquidity before and after the onset of the subprime crisis.” *Journal of Financial Economics*, 103, 471–492.

- Duffie, Darrell (1999), “Credit swap valuation.” *Financial Analysts Journal*, 55, 73–87.
- Düllmann, Klaus, Martin Scheicher, and Christian Schrieder (2007), “Asset correlations and credit portfolio risk: an empirical analysis.” Technical report, Discussion Paper, Series 2: Banking and Financial Supervision.
- Easley, David and Maureen O’Hara (1987), “Price, trade size, and information in securities markets.” *Journal of Financial Economics*, 19, 69–90.
- Eberhart, Allan C, William T Moore, and Rodney L Roenfeldt (1990), “Security pricing and deviations from the absolute priority rule in bankruptcy proceedings.” *The Journal of Finance*, 45, 1457–1469.
- Eom, Young Ho, Jean Helwege, and Jing-zhi Huang (2004), “Structural models of corporate bond pricing: An empirical analysis.” *Review of Financial Studies*, 17, 499–544.
- Fama, Eugene F and Kenneth R French (1993), “Common risk factors in the returns on stocks and bonds.” *Journal of Financial Economics*, 33, 3–56.
- Fan, Hua and Suresh M Sundaresan (2000), “Debt valuation, renegotiation, and optimal dividend policy.” *Review of Financial Studies*, 13, 1057–1099.
- Fei, Fei, Ana-Maria Fuertes, and Elena Kalotychou (2012), “Credit rating migration risk and business cycles.” *Journal of Business Finance & Accounting*, 39, 229–263.
- Feldhütter, Peter (2012), “The same bond at different prices: identifying search frictions and selling pressures.” *Review of Financial Studies*, 25, 1155–1206.
- Feller, William (1951), “Two singular diffusion problems.” *Annals of Mathematics*, 54, 173–182.
- Fernández, Carmen and Mark FJ Steel (1998), “On bayesian modeling of fat tails and skewness.” *Journal of the American Statistical Association*, 93, 359–371.

- Figlewski, Stephen, Halina Frydman, and Weijian Liang (2012), “Modeling the effect of macroeconomic factors on corporate default and credit rating transitions.” *International Review of Economics & Finance*, 21, 87–105.
- Fons, Jerome S (1994), “Using default rates to model the term structure of credit risk.” *Financial Analysts Journal*, 50, 25–32.
- Franks, Julian R and Walter N Torous (1994), “A comparison of financial recontracting in distressed exchanges and chapter 11 reorganizations.” *Journal of Financial Economics*, 35, 349–370.
- Frazzini, Andrea and Lasse Heje Pedersen (2014), “Betting against beta.” *Journal of Financial Economics*, 111, 1–25.
- Frydman, Halina and Til Schuermann (2008), “Credit rating dynamics and markov mixture models.” *Journal of Banking & Finance*, 32, 1062–1075.
- Gebhardt, William R, Soeren Hvidkjaer, and Bhaskaran Swaminathan (2005), “The cross-section of expected corporate bond returns: Betas or characteristics?” *Journal of Financial Economics*, 75, 85–114.
- Geske, Robert (1977), “The valuation of corporate liabilities as compound options.” *Journal of Financial and Quantitative Analysis*, 12, 541–552.
- Glosten, Lawrence R and Lawrence E Harris (1988), “Estimating the components of the bid/ask spread.” *Journal of Financial Economics*, 21, 123–142.
- Glosten, Lawrence R and Paul R Milgrom (1985), “Bid, ask and transaction prices in a specialist market with heterogeneously informed traders.” *Journal of Financial Economics*, 14, 71–100.
- Goh, Jeremy C and Louis H Ederington (1993), “Is a bond rating downgrade bad news, good news, or no news for stockholders?” *The Journal of Finance*, 48, 2001–2008.

- Goldstein, Michael A, Edith S Hotchkiss, and Erik R Sirri (2007), “Transparency and liquidity: A controlled experiment on corporate bonds.” *Review of Financial Studies*, 20, 235–273.
- Güttler, André and Peter Raupach (2008), “The impact of downward rating momentum on credit portfolio risk.” Technical report, Discussion Paper, Series 2: Banking and Financial Supervision.
- Güttler, Andre and Peter Raupach (2010), “The impact of downward rating momentum.” *Journal of Financial Services Research*, 37, 1–23.
- Hamilton, David T and Richard Cantor (2004), “Rating transition and default rates conditioned on outlooks.” *The Journal of Fixed Income*, 14, 54–70.
- Hansen, Bruce E (1994), “Autoregressive conditional density estimation.” *International Economic Review*, 35, 705–730.
- Hanson, Samuel and Til Schuermann (2006), “Confidence intervals for probabilities of default.” *Journal of Banking & Finance*, 30, 2281–2301.
- Hasbrouck, Joel (2009), “Trading costs and returns for us equities: Estimating effective costs from daily data.” *The Journal of Finance*, 64, 1445–1477.
- Hasbrouck, Joel and Duane J Seppi (2001), “Common factors in prices, order flows, and liquidity.” *Journal of Financial Economics*, 59, 383–411.
- Haugen, Robert A and A James Heins (1975), “Risk and the rate of return on financial assets: Some old wine in new bottles.” *Journal of Financial and Quantitative Analysis*, 10, 775–784.
- Hendricks, Darryll (1996), “Evaluation of value-at-risk models using historical data (digest summary).” *Economic Policy Review Federal Reserve Bank of New York*, 2, 39–67.
- Hill, Paula, Robert Brooks, and Robert Faff (2010), “Variations in sovereign credit quality assessments across rating agencies.” *Journal of Banking & Finance*, 34, 1327–1343.

- Holden, Craig W (2009), “New low-frequency spread measures.” *Journal of Financial Markets*, 12, 778–813.
- Hong, Harrison, Terence Lim, and Jeremy C Stein (2000), “Bad news travels slowly: Size, analyst coverage, and the profitability of momentum strategies.” *The Journal of Finance*, 55, 265–295.
- Houweling, Patrick, Albert Mentink, and Ton Vorst (2005), “Comparing possible proxies of corporate bond liquidity.” *Journal of Banking & Finance*, 29, 1331–1358.
- Houweling, Patrick and Jeroen Van Zundert (2014), “Factor investing in the corporate bond market.” *Available at SSRN: 2516322*.
- Huisman, Ronald, Kees G Koedijk, and Rachel AJ Pownall (1998), “Var-x: fat tails in financial risk management.” *Journal of Risk*, 1, 47–61.
- Hull, John C and Alan White (2000), “Valuing credit default swaps ii: modeling default correlations.” *NYU Working Paper No. FIN-00-022*.
- International Swaps and Derivatives Association (2010), “Isda market survey; notional amounts outstanding at year-end, all surveyed contracts, 1987-present.” Technical report, ISDA.
- Israel, Robert B, Jeffrey S Rosenthal, and Jason Z Wei (2001), “Finding generators for markov chains via empirical transition matrices, with applications to credit ratings.” *Mathematical Finance*, 11, 245–265.
- Jacobson, Tor, Jesper Lindé, and Kasper Roszbach (2006), “Internal ratings systems, implied credit risk and the consistency of banks risk classification policies.” *Journal of Banking & Finance*, 30, 1899–1926.
- Jarrow, Robert A, David Lando, and Stuart M Turnbull (1997), “A markov model for the term structure of credit risk spreads.” *Review of Financial studies*, 10, 481–523.

- Jegadeesh, Narasimhan and Sheridan Titman (1993), “Returns to buying winners and selling losers: Implications for stock market efficiency.” *The Journal of Finance*, 48, 65–91.
- Jensen, Michael C (1968), “The performance of mutual funds in the period 1945–1964.” *The Journal of Finance*, 23, 389–416.
- Jensen, Michael C, Fischer Black, and Myron S Scholes (1972), “The capital asset pricing model: Some empirical tests.” *Studies in the Theory of Capital Markets*.
- Johnson, Ben, Horternse Bioy, and Dimitar Boyadzhiev (2016), “Assessing the true cost of strategic-beta etfs.” Technical report, Morningstar.
- Jondeau, Eric, Ser-Huang Poon, and Michael Rockinger (2007), *Financial modeling under non-Gaussian distributions*. Springer Science & Business Media.
- Jones, E Philip, Scott P Mason, and Eric Rosenfeld (1984), “Contingent claims analysis of corporate capital structures: An empirical investigation.” *Journal of Finance*, 611–625.
- Jostova, Gergana, Stanislava Nikolova, Alexander Philipov, and Christof W Stahel (2013), “Momentum in corporate bond returns.” *Review of Financial Studies*, 21.
- Kat, Harry M and Chris Brooks (2001), “The statistical properties of hedge fund index returns and their implications for investors.” *Cass Business School Research Paper*.
- Kaufmann, Roger (2004), “Long-term risk management.” *ETH Zurich: Zurich, Switzerland*.
- Keating, Con and William F Shadwick (2002), “A universal performance measure.” *Journal of Performance Measurement*, 6, 59–84.
- Kerry, Will (2008), “Measuring financial market liquidity.” *Journal of Risk Management in Financial Institutions*, 1, 181–190.
- Khang, Kenneth and Tao-Hsien Dolly King (2004), “Return reversals in the bond market: evidence and causes.” *Journal of Banking & Finance*, 28, 569–593.

- Kidd, Deborah (2014), “Factor investing: When alpha becomes beta.” *Investment Risk and Performance Feature Articles*, 2014.
- Kim, In Joon, Krishna Ramaswamy, and Suresh Sundaresan (1993), “Does default risk in coupons affect the valuation of corporate bonds?: A contingent claims model.” *Financial Management*, 117–131.
- Knight, Frank H (1921), “Risk, uncertainty and profit.” *New York: Hart, Schaffner and Marx*.
- Koziol, Christian and Peter Sauerbier (2007), “Valuation of bond illiquidity: An option-theoretical approach.” *The Journal of Fixed Income*, 16, 81.
- Kraft, Holger and Mogens Steffensen (2007), “Bankruptcy, counterparty risk, and contagion.” *Review of Finance*, 11, 209–252.
- Kyle, Albert S (1985), “Continuous auctions and insider trading.” *Econometrica: Journal of the Econometric Society*, 53, 1315–1335.
- Lando, David (2009), *Credit risk modeling: theory and applications*. Princeton University Press.
- Lando, David and Torben M Skødeberg (2002), “Analyzing rating transitions and rating drift with continuous observations.” *Journal of Banking & Finance*, 26, 423–444.
- Leland, Hayne E (1994), “Corporate debt value, bond covenants, and optimal capital structure.” *The Journal of Finance*, 49, 1213–1252.
- Leland, Hayne E and Klaus Bjerre Toft (1996), “Optimal capital structure, endogenous bankruptcy, and the term structure of credit spreads.” *The Journal of Finance*, 51, 987–1019.
- Lesmond, David A, Joseph P Ogden, and Charles A Trzcinka (1999), “A new estimate of transaction costs.” *Review of Financial Studies*, 12, 1113–1141.
- Li, David X (1999), “On default correlation: A copula function approach.” *Available at SSRN 187289*.

- Lin, Ji-Chai, Gary C Sanger, and G Geoffrey Booth (1995), “Trade size and components of the bid-ask spread.” *Review of Financial Studies*, 8, 1153–1183.
- Longstaff, Francis A and Eduardo S Schwartz (1995), “A simple approach to valuing risky fixed and floating rate debt.” *The Journal of Finance*, 50, 789–819.
- Lopez, Jose A (2004), “The empirical relationship between average asset correlation, firm probability of default, and asset size.” *Journal of Financial Intermediation*, 13, 265–283.
- Lyden, S and D Saraniti (2000), “An empirical investigation of the classical theory of corporate security valuation.” *Barclays Global investors*.
- Madhavan, Ananth and Seymour Smidt (1991), “A bayesian model of intraday specialist pricing.” *Journal of Financial Economics*, 30, 99–134.
- Malkiel, Burton G (2007), *A random walk down Wall Street: The time-tested strategy for successful investing*. WW Norton & Company.
- Markit (2008), “Markit iboxx eur/gbp bond price consolidation rules.” Technical report, Markit Group Limited.
- Markit (2012a), “Markit iboxx gbp benchmark index.” Technical report, Markit Group Limited.
- Markit (2012b), “Markit iboxx rating methodology.” Technical report, Markit Group Limited.
- Markit (2014), “Markit iboxx bond index calculus.” Technical report, Markit Group Limited.
- Markowitz, Harry (1952), “Portfolio selection.” *The Journal of Finance*, 7, 77–91.
- Markowitz, Harry M (1991), “Foundations of portfolio theory.” *The Journal of Finance*, 46, 469–477.
- Mason, Scott P and Sudipto Bhattacharya (1981), “Risky debt, jump processes, and safety covenants.” *Journal of Financial Economics*, 9, 281–307.

- Merton, Robert C (1974), “On the pricing of corporate debt: The risk structure of interest rates*.” *The Journal of Finance*, 29, 449–470.
- Merton, Robert C (1987), “A simple model of capital market equilibrium with incomplete information.” *The Journal of Finance*, 42, 483–510.
- Moody’s Investor Services (2011), “Corporate default and recovery rates, 1920–2010.” Technical report, Moody’s Investor Services.
- Mueller, Philippe (2008), “Credit spreads and real activity.” In *Workshop, Bank for International Settlements*, 30.
- Nickell, Pamela, William Perraudin, and Simone Varotto (2000), “Stability of rating transitions.” *Journal of Banking & Finance*, 24, 203–227.
- Niederhoffer, Victor and Matthew Fontaine Maury Osborne (1966), “Market making and reversal on the stock exchange.” *Journal of the American Statistical Association*, 61, 897–916.
- Nielsen, Lars Tyge, Jesus Saà-Requejo, and Pedro Santa-Clara (2001), *Default risk and interest rate risk: The term structure of default spreads*. INSEAD.
- Ogden, Joseph P (1987), “Determinants of the ratings and yields on corporate bonds: Tests of the contingent claims model.” *Journal of Financial Research*, 10, 329–340.
- O’Hara, Maureen (1995), *Market microstructure theory*, volume 108. Blackwell Cambridge, MA.
- Poor’s, Standard & (2014), “Understanding credit ratings.” Technical report, Standard & Poor’s.
- Pritchard, DJ (2006), “Modeling disability in long-term care insurance.” *North American Actuarial Journal*, 10, 48–75.
- Rachev, Svetlozar T, Christian Menn, and Frank J Fabozzi (2005), *Fat-tailed and skewed asset return distributions: implications for risk management, portfolio selection, and option pricing*, volume 139. John Wiley & Sons.

- Reinganum, Marc R (1981), “Misspecification of capital asset pricing: Empirical anomalies based on earnings’ yields and market values.” *Journal of Financial Economics*, 9, 19–46.
- Roll, Richard (1984), “A simple implicit measure of the effective bid-ask spread in an efficient market.” *The Journal of Finance*, 39, 1127–1139.
- Salmon, Felix (2009), “Recipe for Disaster: The Formula That Killed Wall Street.” <http://http://www.wired.com/2009/02/wp-quant/>.
- Sarig, Oded and Arthur Warga (1989), “Some empirical estimates of the risk structure of interest rates.” *The Journal of Finance*, 44, 1351–1360.
- Sefton, James, David Jessop, Giuliano De Rossi, Claire Jones, and Heran Zhang (2011), “Low-risk investing.” *UBS Investment Research*.
- Sengupta, Ambar (2005), *Pricing derivatives: The financial concepts underlying the mathematics of pricing derivatives*. McGraw-Hill.
- Service, Moody’s Investors (2014), “Annual default study: Corporate default and recovery rates, 1920-2013.” Technical report, Moody’s.
- Shapiro, Samuel Sanford and Martin B Wilk (1965), “An analysis of variance test for normality (complete samples).” *Biometrika*, 52, 591–611.
- Sharpe, William F (1964), “Capital asset prices: A theory of market equilibrium under conditions of risk.” *The Journal of Finance*, 19, 425–442.
- Shimko, David C, Naohiko Tejima, and Donald R Van Deventer (1993), “The pricing of risky debt when interest rates are stochastic.” *The Journal of Fixed Income*, 3, 58–65.
- Sloan, R (1996), “Do stock prices fully reflect information in accruals and cash flows about future earnings?” *Accounting Review*, 71, 289–315.
- Stoll, Hans R (1989), “Inferring the components of the bid-ask spread: theory and empirical tests.” *The Journal of Finance*, 44, 115–134.

- Stulz, René M (2009), “Credit default swaps and the credit crisis.” Technical report, National Bureau of Economic Research.
- Tang, Dragon Yongjun and Hong Yan (2010), “Market conditions, default risk and credit spreads.” *Journal of Banking & Finance*, 34, 743–753.
- Van Dijk, Mathijs A (2011), “Is size dead? a review of the size effect in equity returns.” *Journal of Banking & Finance*, 35, 3263–3274.
- van Loon, P.R.F., A.J.G. Cairns, A.J. McNeil, and A. Veys (2015), “Modelling the liquidity premium on corporate bonds.” *Annals of Actuarial Science*, 9, 264–289.
- Vasicek, Oldrich (1977), “An equilibrium characterization of the term structure.” *Journal of Financial Economics*, 5, 177–188.
- Vassalou, Maria and Yuhang Xing (2004), “Default risk in equity returns.” *The Journal of Finance*, 59, 831–868.
- Vazza, Diane and Nick W Kraemer (2014), “Default, transition, and recovery: 2013 annual global corporate default study and rating transitions.” Technical report, Standards & Poor’s Ratings Services.
- Webber, Lewis (2007), “Decomposing corporate bond spreads.” *Bank of England Quarterly Bulletin*.
- Xing, Haipeng, Ning Sun, and Ying Chen (2012), “Credit rating dynamics in the presence of unknown structural breaks.” *Journal of Banking & Finance*, 36, 78–89.
- Zeng, Bin and Jing Zhang (2001), “An empirical assessment of asset correlation models.” *Moody’s KMV Research Paper*.
- Zhang, X (2006), “Information uncertainty and stock returns.” *The Journal of Finance*, 61, 105–137.
- Zhou, Chunsheng (2001), “The term structure of credit spreads with jump risk.” *Journal of Banking & Finance*, 25, 2015–2040.